

Biased Promotions

Finance Working Paper N° 1094/2025

June 3, 2026

Daniel Ferreira

London School of Economics, CEPR and ECGI

Radoslaw Nikolowa

Queen Mary University of London

Elena Pikulina

University of British Columbia

© Daniel Ferreira, Radoslaw Nikolowa and Elena Pikulina 2026. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

This paper can be downloaded without charge from:
<https://www.ecgi.global/publications/working-papers>

ECGI Working Paper Series in Finance

Biased Promotions

Working Paper N° 1094/2025
June 3, 2026

Daniel Ferreira
Radoslaw Nikolowa
Elena Pikulina

We thank Ashwini Agrawal, Sonia Falconieri, Eric French, Rachel Kranton, Louis-Pierre Lepage, Lubomir Litov, Michael Powell, Michael Waldman, David Webb, Adam Winegar, seminar participants at LSE, QMUL, São Paulo School of Economics (FGV), Cambridge University, HEC Montréal, University of Mannheim, Goethe University Frankfurt, Frankfurt School of F?

© Daniel Ferreira, Radoslaw Nikolowa and Elena Pikulina 2026. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Abstract

We present a model of biased promotions: workers differ only by a nonproductive label, "Blue" or "Red," and firms favor Blue workers in promotion decisions. In equilibrium, worker self-sorting implies (partial) segregation and endogenous firm heterogeneity. Large, high-wage firms offer risky career paths, attracting workers from both groups, whereas small, low-wage firms offer stable careers that attract only Red workers. Promotion biases can benefit firms by weakening workers' outside options and increasing industry profits. The model generates persistent group differences in promotions, earnings, and career trajectories as an equilibrium outcome of competitive labor markets.

Keywords:

Daniel Ferreira

London School of Economics, CEPR and ECGI
Professor of Finance
London School of Economics
e-mail: d.ferreira@lse.ac.uk

Radoslaw Nikolowa

Queen Mary University of London
Reader
e-mail: r.nikolowa@qmul.ac.uk

Elena Pikulina

University of British Columbia
Assistant Professor
e-mail: elena.pikulina@sauder.ubc.ca

Biased Promotions*

Daniel Ferreira

Radoslaw Nikolowa

Elena S. Pikulina

June 3, 2026

Abstract

We present a model of biased promotions: workers differ only by a nonproductive label, “Blue” or “Red,” and firms favor Blue workers in promotion decisions. In equilibrium, worker self-sorting implies (partial) segregation and endogenous firm heterogeneity. Large, high-wage firms offer risky career paths, attracting workers from both groups, whereas small, low-wage firms offer stable careers that attract only Red workers. Promotion biases can benefit firms by weakening workers’ outside options and increasing industry profits. The model generates persistent group differences in promotions, earnings, and career trajectories as an equilibrium outcome of competitive labor markets.

Keywords: Promotions, Discrimination, Inequality, Wages, Careers

*We thank Ashwini Agrawal, Sonia Falconieri, Eric French, Rachel Kranton, Louis-Pierre Lepage, Lubomir Litov, Michael Powell, Michael Waldman, David Webb, Adam Winegar, seminar participants at LSE, QMUL, São Paulo School of Economics (FGV), Cambridge University, HEC Montréal, University of Mannheim, Goethe University Frankfurt, Frankfurt School of Finance and Management, University of Calgary, and participants at The Asian Conference on Organizational Economics at HKU, the BSE Summer Forum Workshop on Tournaments, the Venice Workshop on Labor and Finance, the RAPS/RCFS Europe Meeting, the 2025 FMA Annual Meeting, the NBER Organizational Economics Meeting, the Queensland Corporate Finance Conference, the CEPR Paris Symposium, the conference in honor of Ghazala Azmat, the 2026 UK Women in Finance Conference, the BFI/University of Chicago 2026 Conference on Discrimination, and the Lake District Conference in Corporate Finance. An earlier version circulated under the title “Promotion-Driven Inequality.” The authors have not received specific funding for this research. Ferreira: London School of Economics, CEPR and ECGI; d.ferreira@lse.ac.uk. Nikolowa: Queen Mary University of London; r.nikolowa@qmul.ac.uk. Pikulina: elena.pikulina@sauder.ubc.ca, University of British Columbia.

“There was a time when women with the highest levels of education were barred in subtle and more obvious ways from many high-prestige and high-income occupations and were hired in only a small number of female-dominated occupations, such as teacher, nurse, librarian, and social worker.”

Goldin (2014)

1 Introduction

While occupational segregation by gender has narrowed, significant gender promotion gaps persist in high-skilled occupations. Promotion gaps have been documented among central bankers (Hospido, Laeven, and Lamo (2019)), academic economists (Bosquet et al. (2019)), board directors (Cziraki and Robertson (2021)), lawyers (Azmat, Cuñat, and Henry (2025)), and bankers (Bircan, Friebe, and Stahl (2024); Huang, Mayer, and Miller (2024); Ceccarelli, Herpfer, and Ongena (2024)), among others. Much of the literature attributes these gaps to biased promotion decisions. For example, Benson, Li, and Shue (2026) find that women in management-track careers have better pre-promotion performance than men, yet men are promoted more often based on subjective assessments of greater “potential.” Minni, Nguyen, and Sarsons (2025) find that managers from countries with more progressive gender attitudes promote women at higher rates. Cullen and Perez-Truglia (2023) show that social interactions with managers can explain about a third of the gender gap in promotions. Although promotion gaps are pervasive and economically consequential, accounting for a substantial share of the gender earnings gap,¹ their broader labor market consequences remain largely unexplored.

In this paper, we examine how biased promotion decisions shape labor market equilibrium outcomes. Specifically, we ask: How do these biases affect firm-level wage structures? Do they affect labor demand, firm size, and economic efficiency? Do they distort labor supply decisions, in particular, the career paths workers pursue? Do they contribute to labor market segmentation? Answering these questions requires a theoretical framework that captures how promotion biases influence both firm behavior and workers’ career paths.

To address these questions, we present a model of promotions in which workforce size, wages, and internal career paths are determined in a competitive labor market equilibrium. A key feature of our model is that top positions carry greater prestige. As Baker, Jensen, and Murphy (1988) note, employees “*value the pay and prestige associated with a higher rank in the organization*” (p. 599, emphasis added). We take this idea seriously by explicitly modeling employees’ preferences over pay and prestige, as in Ferreira and Nikolowa (2024). In the model, some jobs are perceived

¹Cullen and Perez-Truglia (2023) speak of “*a consensus that this gap is primarily due to differences in promotion rates*” (p. 1708). Recent evidence in Arellano-Bover et al. (2024) is also consistent with this view.

as being of higher quality than others. We can think of “high-quality jobs” as more prestigious or offering other valuable amenities. The equilibrium features compensating differentials à la Rosen (1986) in wages and promotion probabilities (see Mas (2025) for a comprehensive review of compensating differentials models).

The model is as follows. There is a large number of *ex ante* identical firms. Each firm has a limited number of high-quality jobs. Workers live for two periods, and have the same skills. Firms can create several entry-level positions in a “low-quality job,” fill them with young workers, and then promote a subset of those workers to high-quality jobs. Because all workers are equally qualified for the top job, workers perceive the promotion decision as a (possibly biased) lottery. In practice, promotion uncertainty may operate through uneven opportunities to build promotable skills: access to promotable tasks, high-visibility projects, or mentoring. It may also stem from within-firm contests that are meritocratic yet uncertain in outcomes. What matters is that workers view promotions as *ex ante* random. We call this mechanism a *promotion lottery*.

We assume that workers differ only by a payoff-irrelevant label, “Blue” or “Red,” yet firms favor Blue workers when choosing among equally qualified candidates for promotion. This form of discrimination is “subtle”: it cannot be objectively verified and has no direct payoff consequences to firms (as in Pikulina and Ferreira (2026)). We interpret subtle discrimination as a form of friction: firms are unable to commit to promoting Blue and Red workers with equal probability. In practice, subtle discrimination may take the form of Red workers receiving fewer promotable assignments (Babcock et al. (2017); Bircan et al. (2024)) or having less access to mentoring (Athey et al. (2000); Matsa and Miller (2011); Kunze and Miller (2017)). We treat such biases as exogenously given and abstract from belief updating (Bohren et al. (2019)).

Partly due to the threat of legal action, overt discrimination has become relatively rare, making subtle discrimination a more relevant phenomenon. For example, subtle biases could manifest themselves in the form of gender norms, which have been shown to impact gender differences in labor market outcomes (Cortés et al. (2026)). Bennedsen et al. (2022) show that transparency about gender gaps increases women’s promotion probabilities, suggesting a role for subtle biases in explaining gender promotion gaps. Consistent with our modeling of subtle discrimination, Huber, Lindenthal, and Waldinger (2021) and Ronchi and Smith (2026) show that firms may engage in discrimination without immediate profit consequences when workers are closely substitutable in terms of their skills.

When firms are strongly biased toward Blue workers, the equilibrium exhibits firm heterogeneity and sorting. Two types of firms endogenously emerge. *Mixed firms* hire both Blue and Red workers. These firms create internal labor markets with *risky career paths*: they hire several entry-level workers, but promote only some of them to the top-level positions. Blue workers constitute the majority in Mixed firms and have a higher probability of promotion than Red workers (i.e.,

there is a Blue-Red promotion gap). In contrast, *Red firms* attract only Red workers. They offer *safe career paths* in which employees are promoted with certainty.

Red firms are smaller and pay lower entry-level wages than Mixed firms. Thus, in equilibrium, high- and low-wage firms coexist. Moreover, the unfavored group (Red) is less likely to work in high-wage firms, and even when they do, their expected wages are lower due to biased promotion decisions. Our model therefore offers a new explanation for the evidence that women are underrepresented in high-wage firms, and earn less than men in such firms (Bertrand and Hallock (2001); Card, Cardoso, and Kline (2016); Card et al. (2018); Goldin et al. (2017); Huneus et al. (2021)). Crucially, we can rationalize these findings in a model in which firm heterogeneity arises endogenously and is not driven by differences in workers' preferences. Comparative statics reveal that wages in Red firms decrease with the promotion bias, while wages in Mixed firms *increase* with the bias. Consequently, the wage gap between the two types of firms widens as the bias grows.

Our analysis reveals a striking result: when the bias is large, *profits increase with the bias*. The mechanism is that a larger bias weakens workers' bargaining position, giving firms the upper hand in the labor market: the bias functions as an implicit collusion device, depressing the outside options of Red workers and allowing firms to capture a larger share of the surplus. A direct consequence is that the favored group (Blue) may paradoxically be harmed by a bias in its favor. Unexpectedly, our model offers a potential rationalization of some postmodern perspectives that argue that discrimination in the labor market causes worker fragmentation, potentially limiting collective bargaining power. This outcome emerges under perfect competition and rational expectations.

In our model, competition does not undo discrimination. Because subtle biases are unobservable at the firm level, firm entry cannot drive discriminating firms out of the market. A larger bias typically induces more firms to enter. Entry may increase Red-Blue inequality and, in some cases, reduce the utilities of all workers. If the marginal entrant is an inefficiently small Red firm, total surplus falls with entry, as more workers are employed by firms with suboptimal organizational structures. With relatively more labor demand coming from inefficiently small firms, workers' expected earnings fall.

Our paper is related to the literature on promotions and internal labor markets (Doeringer and Piore (1971); Baker, Gibbs, and Holmstrom (1994)). Huitfeldt et al. (2023) provide large-scale evidence of internal labor markets across several industries in Norway. They find evidence that firms have multiple internal labor markets with well-defined ports of entry, insider bias in promotions, and a strong correlation between rank and wages. Using a similar methodology, Ewens and Giroud (2024) find evidence of internal labor markets in large US firms, characterized by pyramidal hierarchies and slot constraints on promotions. See also Lazear, Shaw, and Stanton (2018) for a theory and evidence on the importance of competition for job slots.

The theoretical literature on promotions and internal labor markets is extensive. Classic works include Lazear and Rosen (1981), Waldman (1984), Prendergast (1993), and Gibbons and Waldman (1999), while more recent contributions include dynamic moral hazard models by Axelson and Bond (2015) and Ke, Li, and Powell (2018). Most of this work emphasizes moral hazard, selection, or matching as explanations for promotion practices.

Fahn and Klein (2025) emphasize the need for alternative theories of promotions, especially in light of recent empirical puzzles, such as the evidence on the “Peter Principle” (Benson, Li, and Shue (2019)). Similar to our work, they present a model in which firms design promotion practices to extract workers’ surpluses optimally. Their model is based on the exploitation of workers’ overconfidence. In contrast, in this paper (as in Ferreira and Nikolowa (2024, 2025)), firms design contracts to optimally exploit workers’ combined preferences for wages and job amenities.

There is abundant evidence of managerial biases in personnel decisions other than promotions. Hoffman, Kahn, and Li (2018) show that biases, rather than superior information, account for the majority of instances in which managers exercise discretion to select candidates. Using a comprehensive audit study, Kline, Rose, and Walters (2022) document evidence of subtle forms of racial discrimination in hiring by large firms. While audit studies can precisely estimate average firm-specific biases, unique instances of subtle discrimination remain essentially undetectable. However, it is unclear whether audit studies can be used to detect biases in internal promotions.

This paper is also related to the extensive literature on discrimination in labor markets (see Fang and Moro (2011), Lang and Lehmann (2012), and Onuchic (2023) for reviews). Closely related is the seminal work of Lazear and Rosen (1990), who provide the first theoretical analysis of promotion gaps, in a model where men and women have exogenously different outside options. In contrast, in our model, Blue and Red workers have *endogenously* different outside options in equilibrium *because* of discrimination. Finally, our focus on the aggregate consequences of small biases is related to theoretical works on bias amplification, such as Lang, Manove, and Dickens (2005), Bartoš et al. (2016), Davies, Van Wesep, and Waters (2024), Siniscalchi and Veronesi (2021), and Onuchic and Ray (2023), among others.

In Section 2, we present a simplified version of the model. Section 3 presents the full model. Section 4 concludes. All proofs are provided in the Appendix. The Internet Appendix provides analyses of omitted cases and extensions.

2 A Simple Model of Biased Promotions

In this section, we present a simplified version of the model. We assume that there is only one firm, with a fixed size. As a consequence of this assumption, the model in this section cannot generate our main results. Instead, the goal of this section is to explain the economic forces and the model’s

underlying logic. We present the full version of the model in Section 3, where we allow a large number of firms to compete for workers, with no exogenous constraints on the number of workers they can hire.

For readability, we keep the presentation informal in this section and discuss all proofs in the text. We postpone a more rigorous and formal analysis to Section 3.

2.1 A Simple Theory of Promotions

Consider an economy with two identical workers who live for two periods. There are two jobs in this economy. Each job $j \in \{s, h\}$ has a quality attribute θ_j , capturing status, responsibility, autonomy, working conditions, or other positive nonpecuniary amenities. Workers care about job quality and wages: $u(\theta_j, w_j)$. For simplicity, we set $u(\theta_j, w_j) = \theta_j u(w_j)$, with $\theta_j > 0$, $u(w_j) > 0$, $u'(w_j) > 0$ and $u''(w_j) < 0$ for all $w_j > 0$, and $u'(0) = \infty$.² Note that the marginal utility of income increases with the nonpecuniary attribute, which is a standard assumption in the literature on status (see, e.g., Becker, Murphy, and Werning (2005)). This property is crucial for our main results.

The first job, the *standard job*, is in excess supply, has attribute $\theta_s = 1$, and pays a fixed wage $\underline{w}$ per period. We may think of this “job” as self-employment. The second job, the *high-quality job*, is in short supply; there is a single firm that can offer one job slot per period. This job has attribute $\theta_h = \theta > 1$.

Suppose the firm lives forever and, at each period, a new cohort of two young workers enters the labor market, replacing the two workers who retire. The firm can hire only one worker per period, in which case its per-period profit is $\pi = R_h - w_h$, where $R_h > 0$ is the revenue and w_h is the wage. Suppose the firm sets w_h at the beginning of a period, and workers apply for the position. If the firm chooses w_h to fill the vacancy at minimum cost, the wage is determined by

$$\theta u(w_h^*) = u(\underline{w}). \quad (1)$$

If more than one worker applies, the firm chooses one at random. We call this the *spot-wage contract* equilibrium. We may think of the selected worker as someone who is “promoted” from the standard job to the high-quality job.

We contend that a theory of promotions must explain at least four basic facts: (i) wages increase upon promotion; (ii) workers strictly prefer to be promoted (i.e., there are “promotion rents”); (iii) slot constraints imply rationing of “high-quality” jobs; and (iv) firms create internal labor markets. This spot-contract theory fails to account for all four required facts: wages decrease upon promotion ($w_h^* < \underline{w}$), promoted workers receive no rents, and there is no rationing or internal labor markets.

²Our results also go through under more general utility specifications $u(\theta, w)$, provided that $u_\theta > 0$, $u_w > 0$, $u_{ww} < 0$ and $u_{\theta w} > 0$ (see Ferreira and Nikolowa (2024)).

There is an additional problem with this theory: The equilibrium implied by (1) is inefficient. To see this, notice that a worker's expected per-period utility is

$$\frac{1}{2}\theta u(w_h^*) + \frac{1}{2}u(\underline{w}) = u(\underline{w}). \quad (2)$$

Because $\theta u'(w_h^*) > u'(\underline{w})$, if a (hypothetical) social planner changes wages to $w_h^* + \Delta$ and $\underline{w} - \Delta$ for $\Delta > 0$ small, the aggregate wage bill is unchanged, and the workers' expected utility becomes

$$\frac{1}{2}\theta u(w_h^* + \Delta) + \frac{1}{2}u(\underline{w} - \Delta) > u(\underline{w}). \quad (3)$$

Thus, this is a Pareto improvement, implying that choosing w_h^* is socially inefficient. The source of inefficiency is that *marginal utilities are not equalized* when (1) holds. Note that this happens because $\theta > 1$ (i.e., the high-quality job is more desirable than the standard job).

The possibility of a Pareto improvement implies that the firm does not maximize its profit by choosing w_h^* as in (1). To consider an alternative solution, suppose the firm sells “lottery tickets” to both workers for a fee. The firm hires the winner of the (fair) lottery for a wage $\hat{w}_h$, and the loser works in the standard job. Without loss of generality, assume that only the loser pays a fee, f (equivalently, we can interpret $\hat{w}_h$ as the “winning” wage net of the fee). The workers are willing to enter the lottery if

$$\frac{1}{2}\theta u(\hat{w}_h) + \frac{1}{2}u(\underline{w} - f) \geq u(\underline{w}). \quad (4)$$

The firm chooses $(f, \hat{w}_h)$ to maximize its profit, $\hat{\pi} = R_h - \hat{w}_h + f$, subject to the participation constraint in (4). The first-order conditions imply

$$u'(\underline{w} - f^*) = \theta u'(\hat{w}_h^*). \quad (5)$$

The solution $(\hat{w}_h^*, f^*)$ must satisfy (4) (with equality) and (5). We call this the *spot-lottery contract* equilibrium. Because the firm may choose $f = 0$ if it wishes, it must do at least as well under spot lotteries as under spot wages. Since setting $f = 0$ in (5) implies $\hat{w}_h > \underline{w}$, the participation constraint is slack at this wage, and the firm can increase its profit by charging a strictly positive fee. Thus, we must have $f^* > 0$.

At first blush, it might seem puzzling that risk-averse workers are willing to pay for lotteries with negative net expected values. The reason for this behavior is that the attribute $\theta > 1$ effectively makes workers behave *as if* they are risk lovers. Formally, utilities are state-dependent, and the marginal utility is higher for all wage levels in state h (i.e., when the worker is assigned to the high-quality job) than in self-employment. The spot-lottery contract exploits this convexity, thereby increasing the firm's profits. Crucially important is the assumption that job assignments affect the marginal utility of income.³

³Auriol and Renault (2008), Auriol, Friebe, and Von Bieberstein (2016), and Ferreira and Nikolowa (2024) make this assumption in models of promotions. The literature on status similarly assumes that status increases the marginal

The spot-lottery contract resembles a promotion tournament. Crucially, wages increase upon “promotion” ($\hat{w}_h^* > \underline{w}$), promoted workers enjoy rents ($\theta u(\hat{w}_h^*) > u(\underline{w})$), and high-quality jobs are rationed. However, we still lack an internal labor market.

The practical relevance of spot-lottery contracts is questionable. While we often see modest application fees for positions, it is unlikely that such fees are sufficient to compensate firms for raising wages above the outside market level. The primary practical challenge in implementing spot lotteries is that firms may struggle to collect substantial fees from liquidity-constrained applicants. In our simple model, this could arise because the outside wage $\underline{w}$ cannot be pledged, either because of legal constraints (e.g., the worker would need to pledge a significant share of her future income to pay fees for unsuccessful job applications) or because $\underline{w}$ is not a wage but the money-equivalent of a bundle of pecuniary and nonpecuniary benefits of self-employment.

If firms cannot charge substantial application fees, they may still use their productive technologies to extract more surplus from workers who value the high-quality job. Suppose that the firm may create (at no cost) a new job, which we call job l . We assume that this job has attribute $\theta_l = 1$ (as the standard job). The firm can create as many positions as it wants for this job. Each filled job- l position yields revenue R_l , but we assume $u(R_l) < u(\underline{w})$, so these jobs are individually inefficient. This implies that the firm does not create job- l positions under either spot-wage contracts or spot-lottery contracts. However, the firm now has the option of offering the following long-term employment contract to each young worker: “You work in job l for wage $\tilde{w}_l$ while young. When old, with probability 0.5 you will be promoted to job h and earn wage $\tilde{w}_h$.” We call this contract a *promotion-lottery contract*. This contract creates an internal labor market for the two workers, with an “up-or-out” structure; the worker who is not promoted will prefer to leave the firm and earn $\underline{w}$ in the standard job.

The workers will accept the promotion lottery contract if (assuming no discounting)

$$u(\tilde{w}_l) + \frac{1}{2}\theta u(\tilde{w}_h) + \frac{1}{2}u(\underline{w}) \geq 2u(\underline{w}). \quad (6)$$

The firm chooses $(\tilde{w}_l, \tilde{w}_h)$ to maximize its profit, $\tilde{\pi} = R_h - \tilde{w}_h + 2(R_l - \tilde{w}_l)$, subject to the participation constraint in (6). The first-order conditions imply

$$u'(\tilde{w}_l^*) = \theta u'(\tilde{w}_h^*). \quad (7)$$

This condition implies that *wages must increase upon promotion*. A promotion premium exists because the marginal utility of money is higher at the top job. Tournament models also predict a promotion premium, but for different reasons. In classic tournament models in the tradition of Lazear and Rosen (1981), promotion premiums are used to incentivize workers to exert non-

utility of income; see Hopkins and Kornienko (2004), Becker, Murphy, and Werning (2005), and Ray and Robson (2012), among others.

contractible effort. In market-based tournaments (Waldman (2013)), promotion premiums exist to retain high-skilled workers who have been promoted. By contrast, in promotion-lottery models, workers are (productively) identical, and there is no non-contractible effort. Promotion premiums exist because cost-minimizing contracts must equalize workers' marginal utilities across jobs. Intuitively, if marginal utilities differ across jobs, workers are willing to pay a price to transfer wages from the job with lower marginal utility to the job with higher marginal utility.

Contract $(\tilde{w}_l^*, \tilde{w}_h^*)$ is still inefficient relative to $(f^*, \hat{w}_h^*)$ because working in job l is an inefficient way to pay for the promotion lottery. However, if charging fees is infeasible, the promotion lottery is the second-best contract, as long as the profit under promotion lotteries is larger than the profit under spot wages: if $\tilde{\pi}^* \geq R_h - w_h^* \Leftrightarrow \tilde{w}_h^* - w_h^* \leq 2(R_l - \tilde{w}_l^*)$.

Promotion-lottery contracts exhibit several key properties of internal labor markets observed in the real world, including a positive correlation between pay and rank, insider bias in promotions, rationing of top-level jobs, and the presence of ports of entry, all without invoking incentive issues, asymmetric information, or learning. Of course, a promotion-lottery contract can also be used when some of these other features are present. For example, suppose that before being assigned to job h , a worker must first acquire “managerial skills” by performing a task; call it the “promotable task.” Suppose that only one worker in job l can be assigned to this task, and that in the second period the firm always promotes the worker assigned to the promotable task. Thus, from the workers' perspective, the employment contract is a lottery because before joining the firm, workers do not know who the firm will assign to the promotable task.

For concreteness, from now on, we assume that promotions are strictly meritocratic: the firm always promotes the worker with managerial skills. That is, there are no “ties” when competing for promotion. However, an element of randomness remains if the firm hires two (or more) workers for job l because workers do not know who will be assigned to the promotable task. Our simple model shows that the firm strictly prefers not to commit: profit is maximized exactly when the assignment of the promotable task is (perceived as) random.

2.2 Biased Promotions

Suppose now that the two workers have different labels, $i \in \{b, r\}$, for Blue and Red. As in Pikulina and Ferreira (2026), labels are observable and productively irrelevant. The firm offers a promotion-lottery contract to each worker. Let p_i denote the probability that type i is assigned to the promotable task after joining the firm. Since that assignment determines promotion deterministically, we also refer to p_i as the *promotion probability*.

The firm is biased towards Blue: the promotion probabilities are $p_b = \frac{1+\beta_b}{2}$ and $p_r = \frac{1-\beta_b}{2}$ for Blue and Red, respectively, where $\beta_b \in [0, 1]$. Let $p_b - p_r = \beta_b \geq 0$ denote the *subtle bias*

towards Blue. We assume that subtle biases are exogenously given and fixed, that is, we abstract from belief updating (Bohren et al. (2019)). Promoting either Blue or Red does not directly affect the firm's payoff. Technically speaking, the lack of direct payoff consequences implies that subtle discrimination is neither taste-based nor statistical discrimination.⁴

The existence of a subtle bias implies that the firm cannot commit to a fair promotion lottery. In addition, we assume that the firm cannot *overtly discriminate*: wages $(\tilde{w}_l, \tilde{w}_h)$ must be the same for both worker types. Now, the lifetime participation constraints are type-dependent:

$$u(\tilde{w}_l) + \frac{1 + \beta_b}{2} \theta u(\tilde{w}_h) + \frac{1 - \beta_b}{2} u(\underline{w}) \geq 2u(\underline{w}) \quad (8)$$

for the Blue worker, and

$$u(\tilde{w}_l) + \frac{1 - \beta_b}{2} \theta u(\tilde{w}_h) + \frac{1 + \beta_b}{2} u(\underline{w}) \geq 2u(\underline{w}) \quad (9)$$

for the Red worker. To preserve the property the high-quality job remains desirable to Red, assume $(1 - \beta_b)\theta > 1$. Because we must also have $\theta u(\tilde{w}_h) \geq u(\underline{w})$ in equilibrium (otherwise the firm is unable to retain any old worker in job h), we have that (9) implies (8). Thus, the firm chooses $(\tilde{w}_l, \tilde{w}_h)$ to maximize its profit subject to (9).

The first-order conditions imply

$$u'(\tilde{w}_l^*) = (1 - \beta_b)\theta u'(\tilde{w}_h^*). \quad (10)$$

Note that, as before, wages increase upon promotion. However, unlike the unbiased case, marginal utilities are not equalized across jobs. The marginal rate of substitution between wages in each job (i.e., the ratio of marginal utilities) is now $1 - \beta_b$. Intuitively, the bias β_b reduces the expected value of the wage in the top job for the Red worker. The firm chooses its optimal contract to equalize the Red worker's marginal utilities *as perceived by her*.

Condition (10) implies that the equilibrium is inefficient (relative to the second-best, i.e., unbiased promotion lotteries). To see this, note that the firm's equilibrium profit is decreasing in β_b ,⁵ which implies that the firm would prefer to be unbiased. Thus, we can think of β_b as a friction: firms with positive subtle biases are unable to commit to a fair promotion lottery. This friction implies a smaller *promotion premium* $(\tilde{w}_h^* - \tilde{w}_l^*)$ than the second-best. In other words, the bias against Red workers reduces wage inequality between young and old workers within the same firm. But this is a reduction in “good inequality,” in the sense that the second-best requires a larger promotion premium.

⁴Pikulina and Ferreira (2026) rationalize β_b as the limiting case of a model where workers have small observable differences, and a biased decision-maker has private information about the workers' productive abilities. The bias is subtle because, given the near-identical observable worker productivities, the decision-maker may use his private information to justify any subjective decision; i.e., he can plausibly deny being biased.

⁵Differentiating the Lagrangian with respect to β_b yields $-\frac{\lambda^*}{2} [\theta u(\tilde{w}_h^*) - u(\underline{w})] < 0$, where λ is the (positive) multiplier.

We conclude that a subtle bias towards one type of worker can lead to flatter career wages (i.e., smaller promotion premiums), lower profits, and a less efficient employment contract. However, the simple model in this section has one main limitation: there is a single monopsonistic firm. In the next section, we extend the model to allow for multiple firms that compete for workers.

3 Biased Promotions in Equilibrium

In this section, we extend the simple model from the previous section to allow for competition in the labor market. All proofs are in the Appendix. The footnotes provide further discussions on the assumptions.

3.1 Setup: Firms, Workers, and Contracts

There is a mass F of identical, infinitely-lived firms indexed by $\tau \in [0, F]$. We take F as given until Subsection 3.7, where we endogenize F by allowing for firm entry. Time is discrete, and there is no discounting. In each period, a firm may have vacancies in two (indivisible) jobs, h and l . A worker can perform only one job per period. Firms have a fixed supply of slots in job h , which we set to 1. Firms can also employ a mass $n \geq 1$ of workers in job l . There are no slot constraints for job l ; firms can create as many l -job slots as they wish (i.e., n is a choice variable). A firm receives revenue $R_j > 0$ per period for each unit mass of workers employed in job $j \in \{h, l\}$.⁶

In each period, a mass $E > F$ of young agents enters the labor force and retires after two periods. Each agent has an observable, productively irrelevant label $i \in \{b, r\}$ for Blue and Red. The proportion of type- i agents in the population is α_i , with $\alpha_b + \alpha_r = 1$. As in the previous section, agents derive utility $\theta_j u(w_j)$ from money and job attributes. Job l has attribute $\theta_l = 1$ and job h has attribute $\theta_h = \theta > 1$. All agents—regardless of their labels—have the same preferences. We denote the (per period) utility of an unemployed (or self-employed) agent by $\underline{u}$. We assume

Assumption 1. (i) $u(R_l) + \theta u(R_h) > 2\underline{u}$ and (ii) $u(R_l) < \underline{u}$.

Assumption 1(i) is necessary for promotion lotteries to be (constrained) efficient. Assumption 1(ii) implies that firms do not want to offer l jobs as “standalone jobs;” l jobs exist only as part of a promotion lottery leading to job h .⁷ From now on, we normalize the unemployment utility to zero: $\underline{u} = 0$.

⁶The technology exhibits constant returns to labor; revenue from job l increases linearly with n at a constant marginal rate R_l . We chose this specification for simplicity. Our results can also be derived using an alternative model in which the marginal productivity of labor, $R_l(n)$, is decreasing in n .

⁷The “inefficient job” assumption, 1(ii), is needed only because of the constant returns technology. Under a technology with decreasing returns, employing few workers in job l may be efficient, but in equilibrium firms would hire too many workers and the marginal appointment would still be inefficient.

We assume that firms compete in the labor market by offering long-term employment contracts to young workers, which are designed as promotion-lottery contracts. In these contracts, young workers perform job l (the entry-level job) for one period. In the second period of employment, workers who have acquired managerial skills are promoted to job h ; those who fail to be promoted leave the firm.⁸ Each firm has a unit mass of promotable tasks per period. Firms commit to assigning only insiders to promotable tasks but retain discretion over whom to assign to these tasks. Firms attract job applicants by advertising and committing to n , which is the number of vacancies for job l , and a *career wage schedule*, (w_l, w_h) , where w_j is the (per unit mass of workers) wage for job $j \in \{l, h\}$. We call $c = (n, w_l, w_h)$ an *employment contract*.⁹ A firm's *per contract* profit is thus

$$\pi(c) = R_h - w_h + n(R_l - w_l). \quad (11)$$

3.2 Promotion Probabilities

Let p_i denote the belief that a type- i worker has about his/her promotion probability under a given contract. That is, type- i workers expect to be promoted with probability p_i and to be fired when old with probability $1 - p_i$, in which case they become unemployed and enjoy utility $\underline{u} = 0$. If $n = 1$, the contract implies a degenerate promotion lottery: all entry-level workers are assigned to the promotable task and are later promoted to job h . Thus, if $n = 1$, workers expect a *safe career path*: they expect to be promoted with probability $p_i = 1$. If $n > 1$, not all workers can be promoted and thus $p_i < 1$, implying that workers face a *risky career path*.

Belief p_i may depend on the worker's label, i . If $p_b \neq p_r$, workers expect promotions to be biased. To model such biases, we extend Pikulina and Ferreira's (2026) notion of subtle discrimination to the case of n agents. To make the analysis more intuitive, we interpret n (the number of workers) as an integer in this subsection.

Let n_i denote the number of workers of type $i \in \{b, r\}$ hired under contract c , with $n_b + n_r = n$, and $q = \frac{n_b}{n}$ denote the expected proportion of Blue workers who are hired under this contract. Denote the probability that a specific type- i worker is assigned to the promotable task by $p_i(n, q)$. The probability that the firm promotes one of the n_i type- i workers is thus $n_i p_i(n, q)$. We define the firm's bias towards Blue workers as:

$$\beta_b(n, q) = n[qp_b(n, q) - (1 - q)p_r(n, q)] - (2q - 1). \quad (12)$$

The first term on the right-hand side is the probability that the promoted worker is Blue minus the

⁸The "out" part of the up-or-out contract is not strictly necessary and should not be taken literally. It is straightforward to extend the model to allow unpromoted workers to remain at the firm.

⁹To streamline the analysis, we assume that a dismissed worker receives no compensation from the firm. While severance compensation can sometimes be optimal (as shown in Ferreira and Nikolowa (2024)), its availability does not affect the optimality of promotion lotteries or the qualitative features of the equilibrium.

probability that the promoted worker is Red. The second term is the same difference of probabilities if the firm is unbiased.

To pin down the probability of promotion, we take an axiomatic approach and impose three intuitive properties on the probability function.

Assumption 2. $p_i(n, q)$ must satisfy the following axioms:

- Axiom i. Monotonicity:* $p_i(n, q)$ decreases with q .
- Axiom ii. Size Invariance:* $p_i(n, q)$ is homogeneous of degree minus one in n , that is, $p_i(zn, q) = \frac{p_i(n, q)}{z}$ for all $z > 0$.
- Axiom iii. Constant Returns:* Probability changes due to proportion changes are independent of proportion levels, that is, $p_i(n, q_0 + \varepsilon) - p_i(n, q_0)$ is independent of q_0 for any $\varepsilon < 1 - q_0$.

Axiom (i) states that a higher share of Blue workers makes promotion harder for all workers. Axiom (ii) implies that promotion probabilities scale inversely with the number of candidates: in an unbiased promotion lottery, doubling the number of workers halves each worker's promotion probability. We assume that scaling carries over to the biased case, so the bias is invariant to firm size. Axiom (iii) ensures that the firm does not become intrinsically more or less biased as the proportion of Blue workers changes. Together, axioms (ii) and (iii) imply that firms remain equally biased in equilibrium, even when they differ in size and workforce composition. The following proposition shows that Axioms (i) - (iii) are sufficient to determine the function $p_i(n, q)$.

Proposition 1 (Biased promotion probabilities). *Axioms (i) - (iii) imply that the individual promotion probabilities are*

$$p_b(n, q) = \frac{1 + (1 - q)\beta}{n} \text{ and } p_r(n, q) = \frac{1 - q\beta}{n}, \quad (13)$$

where parameter $\beta \in [0, 1]$ measures the intensity of the bias.

Note that $\beta = 2\beta_b(2, 0.5)$. For the remainder of the paper, we assume that workers expect all firms to be equally biased, that is, all firms share the same β .¹⁰

We interpret $p_i(n, q)$ as a friction: firms cannot commit to a fair promotion rule. Note that $p_i(n, q)$ is a belief; the actual promotion probabilities are unobservable and ex post unverifiable, implying that firms cannot contract on promotion probabilities or adopt promotion quotas for Red and Blue (as these amount to overt discrimination).

¹⁰ For $p_b(n, q)$ to be lower than 1 for any β , we need $1 - q \leq n - 1$. In a discrete setting, we must have at least two people, implying $n \geq 2 \Rightarrow p_b(n, q) \leq 1$. However, for convenience, we later work with a continuous n , thus it is technically possible that $n \in (1, 2)$. To streamline the exposition, we will henceforth ignore this case, but discuss it in the proofs where appropriate.

3.3 Equilibrium Conditions

We model the matching process of firms and workers as arising from directed search. First, all firms simultaneously post (and commit to) contracts. Each firm can post a single contract. Then, all young workers apply for jobs. Workers and firms are infinitesimal and, thus, ignore strategic considerations (i.e., search is competitive, in the sense of Wright et al. (2021)).

Suppose firm τ posts contract $c_\tau = (n_\tau, w_{l\tau}, w_{h\tau})$, while the set of contracts posted by all other firms is $C_{-\tau}$. We denote the set of all contracts by $C = \{c_\tau\} \cup C_{-\tau}$ and refer to the “self-employment” contract as $\underline{c}$. If more than n_τ workers apply to the firm, it must randomly choose among them with equal probabilities, regardless of their types. That is, we assume that firms’ hiring decisions are unbiased, unlike their decisions to allocate promotable tasks.¹¹

Let $q(c_\tau, C_{-\tau})$ denote the proportion of Blue workers firm τ expects to hire when it offers contract c_τ and the set of contracts offered by all other firms is $C_{-\tau}$. Function $q(c_\tau, C_{-\tau})$ represents agents’ *beliefs*. We assume that all workers and firms share the same beliefs. A type- i worker’s lifetime utility under c_τ when all other firms offer $C_{-\tau}$ is

$$U_i(c_\tau, C_{-\tau}) := u(w_{l\tau}) + p_i(n_\tau, q(c_\tau, C_{-\tau}))\theta u(w_{h\tau}), \quad (14)$$

where $p_i(n, q)$ is given by (13). Define

$$\bar{U}_i(c_\tau, C_{-\tau}) := \max_{c \in C_{-\tau} \cup \{\underline{c}\}} U_i(c, C_{-\tau} \cup \{c\}). \quad (15)$$

That is, $\bar{U}_i(c_\tau, C_{-\tau})$ is the maximum utility a type- i worker obtains when free to choose any contract in $C_{-\tau}$ or the self-employment contract, $\underline{c}$ (which gives utility $\underline{u} = 0$).

Let $\mu_{i\tau}(c_\tau, C_{-\tau})$ be the mass of workers of type i who apply to firm τ . We can think of $\mu_{i\tau}(c_\tau, C_{-\tau})$ as firm τ ’s *residual labor supply function*, that is, the supply of type- i workers to firm τ given the contracts of all other firms. We impose the following condition on the residual supply functions.

Condition 1 (Labor Supply Functions). *For all $\tau \in [0, F]$ and $i \in \{b, r\}$, the residual labor supply functions $\mu_{i\tau}(c_\tau, C_{-\tau})$ must have the two following properties:*

1. *Symmetry:* If $c_\tau = c_{\tau'}$, then $\mu_{i\tau}(c_\tau, C_{-\tau}) = \mu_{i\tau'}(c_{\tau'}, C_{-\tau'})$.
2. *Perfect Elasticity:*

$$\mu_{i\tau}(c_\tau, C_{-\tau}) = \begin{cases} n_\tau q_i(c_\tau, C_{-\tau}), & \text{if } U_i(c_\tau, C_{-\tau}) \geq \bar{U}_i(c_\tau, C_{-\tau}) \\ 0, & \text{if } U_i(c_\tau, C_{-\tau}) < \bar{U}_i(c_\tau, C_{-\tau}), \end{cases}$$

¹¹This assumption is made to streamline the presentation and is without loss of generality. In a competitive equilibrium, contracts adjust so that there is no rationing of jobs. Thus, each firm receives the same number of applications as the number of vacancies it posts, and subtle biases in hiring have no impact on whom it hires in equilibrium. While hiring biases may still affect off-path behavior, the equilibrium set we consider can be sustained under the same off-path conditions we discuss below.

where $q_b(c_\tau, C_{-\tau}) := q(c_\tau, C_{-\tau})$ and $q_r(c_\tau, C_{-\tau}) := 1 - q(c_\tau, C_{-\tau})$.

Condition 1.1 says that workers perceive two firms offering the same contract as equivalent. In other words, the firm index τ does not affect workers' decisions to apply for vacancies. Condition 1.2 implies that firms are “price takers” in the labor market:¹² they view their residual supply of workers as perfectly elastic. All else being constant, if a firm offers a contract that is weakly better than the competition, it will fill all of its vacancies. Conversely, if a firm offers a contract slightly inferior to the competition's, it attracts no workers.¹³

We define an equilibrium as follows.

Definition 1 (Equilibrium). *An equilibrium is a set of contracts $C^* = \{c_\tau^* : \tau \in [0, F]\}$, beliefs $q(c_\tau, C_{-\tau})$, and symmetric and perfectly elastic residual labor supply functions $\mu_{i\tau}(c_\tau, C_{-\tau})$ such that:*

1. *Firms maximize profit subject to their residual labor supplies: If $c_\tau^* \in C^*$, then $c_\tau^* \in \arg \max_c \pi(c)$ subject to $\mu_{i\tau}(c, C_{-\tau}^*) > 0$ for at least one $i \in \{b, r\}$;*
2. *Supply equals demand: $n_\tau^* = \mu_{b\tau}(c_\tau^*, C_{-\tau}^*) + \mu_{r\tau}(c_\tau^*, C_{-\tau}^*)$, for all $\tau \in [0, F]$;*
3. *Firms and workers have rational expectations: $\mu_{b\tau}(c_\tau^*, C_{-\tau}^*) = q(c_\tau^*, C_{-\tau}^*)[\mu_{b\tau}(c_\tau^*, C_{-\tau}^*) + \mu_{r\tau}(c_\tau^*, C_{-\tau}^*)]$ for all $\tau \in [0, F]$;*
4. *Aggregate labor supply is feasible: $\int_0^F \mu_{i\tau}(c_\tau^*, C_{-\tau}^*) d\tau \leq \alpha_i E$ for $i \in \{b, r\}$.*

Part 1 states that firms choose contracts to maximize profits, subject to workers maximizing utility, assuming all other firms' contracts are given. Part 2 implies that each firm's labor demand equals its residual labor supply. Part 3 requires agents to have rational expectations in equilibrium; that is, beliefs must be consistent with equilibrium play.

Part 4 requires aggregate labor supply to be no greater than the mass of workers in the sector. If Part 4 holds with equality, we have a *(fully) tight labor market equilibrium*: all workers are employed by the firms in the sector. If, for only one type $i \in \{b, r\}$, we have that Part 4 holds with strict inequality, we have a *partially tight labor market equilibrium*, where we may have unemployed workers of one type. Finally, if Part 4 holds with strict inequality for both types, we have a *slack labor market equilibrium*, in which some workers of both types are unemployed. Because only under tight labor markets is there effective competition among firms, from now on, we restrict attention to (partially or fully) tight-labor-market equilibria.

Our equilibrium conditions are insufficient to guarantee the uniqueness of the equilibrium. There are two potential sources of equilibrium multiplicity. First, Definition 1 does not constrain

¹²Or, more precisely, firms take the set of competing contracts, $C_{-\tau}$, as given.

¹³A possible microfoundation for Condition 1.2 is Bertrand competition in the labor market. In our model, firms have an optimal “capacity” (i.e., an optimal number of workers given market conditions); therefore, Bertrand competition does not imply zero profits. This is reminiscent of Kreps and Scheinkman's (1983) classic model of price competition under capacity constraints.

beliefs off the equilibrium path. Second, there might be more than one set of contracts, beliefs, and supply functions that satisfy all conditions in Definition 1. The next condition places restrictions on beliefs off the equilibrium path.

Condition 2 (Off-equilibrium-path beliefs). *Let C^* denote an equilibrium set of contracts and $q_\tau^* := q(c_\tau^*, C_{-\tau}^*)$ for all $\tau \in [0, F]$. For any deviation $c_\tau^d \neq c_\tau^*$ by firm τ , its associated belief must be*

1. *Individually Rational:*

$$q(c_\tau^d, C_{-\tau}^*) = \begin{cases} 1 & \text{if } U_b(c_\tau^d, C_{-\tau}^*) \geq \bar{U}_b(c_\tau^d, C_{-\tau}^*) \text{ and } U_r(c_\tau^d, C_{-\tau}^*) < \bar{U}_r(c_\tau^d, C_{-\tau}^*) \\ 0 & \text{if } U_b(c_\tau^d, C_{-\tau}^*) < \bar{U}_b(c_\tau^d, C_{-\tau}^*) \text{ and } U_r(c_\tau^d, C_{-\tau}^*) \geq \bar{U}_r(c_\tau^d, C_{-\tau}^*) \end{cases}$$

2. *Perfectly Competitive (whenever possible): if q_τ^* is individually rational under $(c_\tau^d, C_{-\tau}^*)$, then $q(c_\tau^d, C_{-\tau}^*) = q_\tau^*$.*

In words, Condition 2.1 requires off-path beliefs to be consistent with workers' *individual rationality* (i.e., participation) constraints. For a deviation contract c_τ^d , if a Blue worker's IR constraint holds but a Red worker's doesn't, everyone expects only Blue workers to apply to contract c_τ^d ; individual rationality thus requires $q(c_\tau^d, C_{-\tau}^*) = 1$. Similarly, if a Red worker's IR constraint holds but a Blue worker's doesn't, everyone expects only Red workers to apply to contract c_τ^d , and thus $q(c_\tau^d, C_{-\tau}^*) = 0$. While it is reasonable to impose individual rationality on beliefs, it is a weak constraint: for any given set of contracts $(c_\tau^d, C_{-\tau}^*)$, there is generally a large set of beliefs $q(c_\tau^d, C_{-\tau}^*)$ that are consistent with individual rationality. In particular, if both IR constraints hold or if both don't hold, individual rationality places no restrictions on $q(c_\tau^d, C_{-\tau}^*)$.

We say that an off-path belief $q(c_\tau^d, C_{-\tau}^*)$ is *perfectly competitive* if it is equal to the equilibrium on-path belief, q_τ^* . In the spirit of perfect competition, Condition 2.2 requires that firms take q_τ^* as given: as long as contracts are compatible with the workers' participation constraints, firms cannot manipulate the proportion of each type of worker that applies to them. We need the "whenever possible" qualifier because, in some cases, setting $q(c_\tau^d, C_{-\tau}^*) = q_\tau^*$ will violate the individual rationality condition (Condition 2.1).¹⁴ In such cases, Condition 2 requires only individual rationality, and beliefs must change after such a deviation.

While the equilibrium is often unique, multiple equilibria may exist for some parameter constellations. Our qualitative results do not depend on which equilibrium firms and workers coordinate on. Let U_i^* denote the equilibrium lifetime utility of a worker of type i . Note that U_i^* is well-defined because competition implies that all workers of the same type must have the same equilibrium utility. For simplicity, when equilibrium multiplicity persists, we assume firms and

¹⁴This happens if and only if, under belief $q(c_\tau^d, C_{-\tau}^*) = q_\tau^*$, we have either $q_\tau^* > 0$ and $U_b(c_\tau^d, C_{-\tau}^*) < \bar{U}_b(c_\tau^d, C_{-\tau}^*)$ and $U_r(c_\tau^d, C_{-\tau}^*) \geq \bar{U}_r(c_\tau^d, C_{-\tau}^*)$, or $q_\tau^* < 1$ and $U_r(c_\tau^d, C_{-\tau}^*) < \bar{U}_r(c_\tau^d, C_{-\tau}^*)$ and $U_b(c_\tau^d, C_{-\tau}^*) \geq \bar{U}_b(c_\tau^d, C_{-\tau}^*)$.

workers coordinate on the equilibrium with the maximum U_r^* (i.e., Red workers' equilibrium utility).

Condition 3 (Equilibrium selection). *If there are multiple equilibria, select the equilibrium with the largest U_r^* .*

Condition 3 implies that Red workers are not harmed by coordination failures. While not necessary for our qualitative results, we impose Condition 3 to clarify that the impact of biased promotions on Red-Blue inequality is not caused by the selection of a “bad equilibrium.”

3.4 Benchmark: Unbiased firms

As a benchmark, we assume there is only one type of worker, say b , so that $\alpha_b = 1$; we temporarily drop the subscript i for notational simplicity. For each contract $c_\tau = (n_\tau, w_{l\tau}, w_{h\tau})$, we must have $q(c_\tau, C_{-\tau}) = 1$, and all workers must have the same probability of promotion, $p(n_\tau, 1) = \frac{1}{n_\tau}$. A worker's lifetime utility under contract c_τ is

$$U(c_\tau) = u(w_{l\tau}) + \frac{1}{n_\tau} \theta u(w_{h\tau}), \quad (16)$$

and a firm's per-contract profit is

$$\pi(c_\tau) = R_h - w_{h\tau} + n_\tau(R_l - w_{l\tau}). \quad (17)$$

Note that R_h is simply a profit shifter, that is, it is a “free parameter,” which we assume to be sufficiently high so that Assumption 1 holds and firms are willing to operate (i.e., profits are non-negative).

Assuming the parameters are such that a tight-labor-market equilibrium exists, the following proposition shows that it is unique and that all firms offer the same contract.

Proposition 2 (Unbiased equilibrium). *There is a unique tight-labor-market equilibrium contract $c^* = (n^*, w_l^*, w_h^*)$, which is offered by all firms.*

Three intuitive conditions determine the equilibrium (see the proof of Proposition 2 for the omitted steps). First, the marginal utilities in each job must equalize:

$$u'(w_l^*) = \theta u'(w_h^*). \quad (18)$$

This condition is the same as in the simple model of the previous section and implies that wages must increase upon promotion.

Second, we have

$$R_l = w_l^* + \frac{U(c^*) - u(w_l^*)}{u'(w_l^*)}. \quad (19)$$

This condition equates the marginal productivity of job l , R_l , to the marginal cost of hiring young workers. An additional worker costs the firm w_l^* plus the monetary equivalent of the worker's expected utility from job h . A worker's expected utility from job h is her lifetime utility minus the utility from job l , $U(c^*) - u(w_l^*)$. This value is converted into dollars by dividing it by the marginal utility of money, $u'(w_l^*)$. Intuitively, if (19) does not hold, firms would like to either increase or reduce the number of job positions they offer. Condition (19) is analogous to the requirement for equalizing wages and marginal productivities in spot labor markets. It implies that, under a promotion-lottery contract, the entry-level wage must be lower than the (period 1) marginal productivity of young workers.

Finally, supply and demand under a tight labor market imply $n^* = E/F$.

The equilibrium of this benchmark model is constrained-efficient. To see this, notice that the assignment of workers to job h must involve some kind of lottery, because the supply of h jobs is $F < E$, and jobs are indivisible. We know from the previous section that, in such a case, the efficient contract involves firms charging workers application fees. However, if firms cannot collect fees from job applicants, the second-best solution is to use promotion lotteries in which workers are initially assigned to job l . Conditional on the use of promotion lotteries, efficiency requires equalizing marginal utilities across the two jobs, as in (18), and equalizing marginal productivity to marginal cost, as in (19).

Suppose now that $\alpha_b < 1$. As a second benchmark, we temporarily assume that firms are allowed to “overtly discriminate,” that is, firms may offer contracts that are exclusive to a type of worker, $i \in \{b, r\}$. Denote such contracts by c_b and c_r . When firms are subtly biased, workers care about their coworkers' identities because they compete for the same promotable tasks. Suppose a firm hires workers of both types. Then, the firm must offer the same wage for the high-level job to both types: $w_{hb} = w_{hr}$; if not, the firm will assign the cheapest type to the promotable task, and the other type will not accept the contract. Under the same top wage, biased firms are more likely to promote Blue workers. Thus, under the same contract, workers—both Blue and Red—prefer to work in firms with fewer Blue workers.

An equilibrium in which some firms have a mixed workforce of Blue and Red workers does not exist. To see this, note that a firm would want to hire both worker types only if they cost the same. However, because Red workers have a lower probability of promotion than Blue workers, they have lower utility. Thus, the firm would strictly prefer to hire only Red workers, who would accept a contract at lower career wages, provided they do not have to compete with Blue workers for promotion.

Formally, let c_i^0 denote a contract offered to type i where $n_i = 0$. That is, a firm that offers c_i^0 overtly discriminates against type i by refusing to hire workers of that type. Let $c_i^* = c^*$ denote the benchmark contract defined in Proposition 2 when offered to type i only. We then have the

following result.

Proposition 3 (Overt-discrimination equilibrium). *If overt discrimination is possible, in a tight-labor-market equilibrium, only two contracts are offered, (c_b^*, c_r^0) and (c_b^0, c_r^*) . The equilibrium is symmetric and displays full segregation: $\alpha_b F$ firms hire only Blue workers and $(1 - \alpha_b)F$ firms hire only Red workers.*

Proposition 3 implies that, as long as firms can overtly discriminate by offering different contracts to Blue and Red workers, subtle biases do not matter. The unique segregated equilibrium is efficient and egalitarian: both Blue and Red workers expect the same wages and have the same expected utility. Thus, the equilibrium contracts in both benchmark cases (Propositions 2 and 3) are the same. We conclude that subtle biases do not distort equilibrium contracts, as long as firms are allowed to discriminate explicitly and become fully segregated.

Overt discrimination is an unappealing feature. What remains unmodeled are the potential externalities from segregation. For example, segregation may reinforce stereotypes and biases, with negative societal consequences. It may also affect the perceived quality of the different jobs, here assumed to share the same θ .¹⁵ With different job qualities, wages would also differ, and inequalities in career opportunities would arise.

For several good reasons, most advanced economies have made overt discrimination illegal. An unintended consequence of this trend is that firms cannot use employment contracts to “de-bias” their promotion practices.

3.5 Equilibrium Characterization

We now characterize the equilibrium set. Both Blue and Red workers have strictly positive mass, $\alpha_i > 0$ for $i \in \{b, r\}$. Contracts take the form $c_\tau = (n_\tau, w_{l\tau}, w_{h\tau})$ (i.e., contractual discrimination is not possible).

We denote an equilibrium by (C^*, μ, q) , where μ and q are shortcuts for equilibrium residual labor supply and belief functions. We say that an equilibrium displays *full segregation* if $q(c_\tau^*, C_{-\tau}^*) \in \{0, 1\}$ for all $\tau \in [0, F]$. That is, in a full segregation equilibrium, there are no firms with a mixed workforce. We say that firms are *ex post homogeneous* if $c_\tau^* = c^*$ for all $\tau \in [0, F]$. That is, firms are ex post homogeneous if they all offer the same contract in equilibrium. The following result shows that there is no full segregation equilibrium with ex post homogeneous firms.

Lemma 1 (No full segregation under homogeneous firms). *There is no equilibrium with ex post homogeneous firms and full segregation.*

¹⁵See Goldin (2014) for a theory in which job prestige and segregation are endogenously determined.

Proposition 3 shows that firms can sustain full segregation by offering economically identical but group-specific contracts. Lemma 1 establishes that this outcome relies on overt discrimination: once firms are restricted to non-discriminatory contracts, full segregation with ex post homogeneous firms is impossible. The intuition is straightforward. With all firms offering the same risky-career-path contract, in a fully segregated configuration, Blue workers would be strictly better off applying to all-Red firms. Full segregation is therefore possible only if firms offer economically distinct contracts that endogenously sort workers.

The next proposition shows that in a (partially or fully) tight-labor-market equilibrium, at most two (ex post) types of firms may coexist.

Proposition 4 (At most two equilibrium contracts). *In an equilibrium with (partially or fully) tight labor markets, a fraction $s^* \in (0, 1]$ of firms offer contract $c_m^* = (n_m^*, w_{lm}^*, w_{hm}^*)$, and all other firms offer contract $c_r^* = (1, w_{lr}^*, w_{hr}^*)$. Blue workers apply to contract c_m^* only.*

Proposition 4 shows that any equilibrium features at most two contracts: a risky-career-path contract and—if it exists—a safe-career-path contract. Intuitively, firm heterogeneity requires distinct contracts that deliver the same equilibrium profit; since the optimal risky-career-path contract is unique, no two such contracts can coexist. Proposition 4 also implies that contract c_r^* attracts only Red workers, while contract c_m^* may attract both worker types. Accordingly, we call firms offering c_r^* *Red firms* and firms offering c_m^* *Mixed firms* (for simplicity, we keep this terminology even when c_m^* attracts only Blue workers).

In the Internet Appendix, we show the full set of conditions that characterize the equilibrium contracts. Here we briefly discuss some of these conditions. For brevity, in what follows we focus on the case where Mixed firms are strictly mixed (i.e., the mass of Red workers in such firms is non-zero).¹⁶

A positive promotion bias ($\beta > 0$) implies that Red workers have lower promotion probabilities than Blue workers. Thus, we must have $U_r^* < U_b^*$, implying that Red workers are cheaper to employ. A Mixed firm must therefore bind the participation constraint of Red workers, yielding the following marginal condition:

$$R_l = w_{lm}^* + \frac{U_r^* - u(w_{lm}^*)}{u'(w_{lm}^*)}. \quad (20)$$

This condition is equivalent to (19), and thus implies that the marginal productivity of Red workers equals their marginal cost to the firm. Note that the bias β does not *directly* affect this condition, but may do so indirectly through U_r^* . We show in the proof of Proposition 2 that w_{lm}^* is decreasing in U_r^* .

¹⁶The case of Mixed firms that attract Blue workers only is covered in the Appendix.

The second marginal condition for profit maximization is

$$u'(w_{lm}^*) = \theta(1 - q^*\beta)u'(w_{hm}^*). \quad (21)$$

As before, wages increase upon promotion.¹⁷ As in (10), marginal utilities are not equalized across jobs: the bias β reduces the expected value of the top job for Red workers.

If all firms are mixed, they must offer the same contract c_m^* , and the optimal number of workers per firm in a fully-tight labor market must be $n_m^* = \frac{E}{F}$. If, instead, contracts c_m^* and c_r^* coexist, they must deliver the same profit, $\pi(c_m^*) = \pi(c_r^*)$. In this case, c_r^* is determined by Red workers' indifference between c_m^* and c_r^* , $u(w_{lr}^*) + \theta u(w_{hr}^*) = U_r^*$, and the marginal condition for Red firms:

$$u'(w_{lr}^*) = \theta u'(w_{hr}^*). \quad (22)$$

In Red firms, wages also increase upon promotion, and Red workers' marginal utilities are equalized across jobs. While Red firms choose wages efficiently given their career structures, the mere existence of these firms indicates resource misallocation because, in the benchmark model, all firms should offer risky career paths. Thus, Red firms are inefficiently small. By the same token, Mixed firms are too large. Intuitively, biased promotions in some firms distort economy-wide career opportunities.

The next proposition shows that, for sufficiently large β , firms must be ex post heterogeneous in equilibrium.

Proposition 5 (Firm heterogeneity). *Suppose an equilibrium with (partially or fully) tight labor markets exists for any $\beta \in [0, 1]$. There exist $\beta_{max} \geq \beta_{min} > 0$ such that:*

1. *For $\beta < \beta_{min}$, all firms offer the same contract $c_m^* = (n_m^*, w_{lm}^*, w_{hm}^*)$ in equilibrium.*
2. *For $\beta > \beta_{max}$, in equilibrium, a fraction $s^* \in (0, 1)$ of firms offer $c_m^* = (n_m^*, w_{lm}^*, w_{hm}^*)$ and all other firms offer $c_r^* = (1, w_{lr}^*, w_{hr}^*)$.*

Proposition 5 shows that ex-post firm heterogeneity arises only if the bias is sufficiently strong. The intuition is as follows. If the bias is zero, all firms must offer the benchmark equilibrium contract c^* as in Proposition 2. In this case, without loss of generality, we can assume that all firms have (identical) mixed workforces. As β increases, by continuity, a single contract must remain the unique solution until $\beta \geq \beta_{min} > 0$.

¹⁷We note that $\theta(1 - q^*\beta) > 1$ in a tight-labor-market equilibrium.




Figure 1: Equilibrium Regions.

Note: Panel A depicts the three equilibrium regions as a function of β and θ , for $R_l = 0.75$, $\alpha_b = 0.5$, and $e = E/F = 2.0$. Panel B depicts the three equilibrium regions as a function of β and e , for $R_l = 0.75$, $\alpha_b = 0.5$, and $\theta = 10.0$.

Figure 1 shows the possible equilibrium types for different parameter configurations, for $u(w) = \ln w$. For this case, we find that $\beta_{min} = \beta_{max} = \bar{\beta}$, with $\bar{\beta}$ strictly lower than 1 for most parameter configurations. Thus, as β increases, the equilibrium eventually features ex-post heterogeneous firms. For β sufficiently high, we observe full segregation: Red workers sort into firms offering safe career paths, while Blue workers choose risky career paths. This pattern may give the impression that Red workers prefer safe, bureaucratic careers, whereas Blue workers prefer risky, competitive ones. However, Blue and Red workers have identical preferences in our model. Red workers avoid risky-career-path contracts because they anticipate lower promotion probabilities than Blue workers.

The resulting sorting may also give the impression that some firms are more favorable to Red workers than others. This is not the case. Even ex-post heterogeneous firms have the same intrinsic bias,¹⁸ but its effect depends on the competitiveness of promotion contests. In safe-career-path firms, where promotions are effectively non-competitive, this bias does not affect outcomes, whereas in risky-career-path firms, it does. Consequently, observed differences in workforce composition may arise even when firms have identical underlying biases.

As the bias decreases, the equilibrium eventually features partial segregation: the contract offered by Mixed firms becomes sufficiently attractive for Red workers, who are indifferent between Red and Mixed firms, while Blue workers continue to apply only to Mixed firms.¹⁹ The following

¹⁸Specifically, axioms (2.ii) and (2.iii) ensure that firm bias is invariant to firm size and workforce composition.

¹⁹For some values of β , partial and full segregation equilibria co-exist, so we apply the selection criterion described

proposition shows that Mixed firms pay their young workers more than Red firms.

Proposition 6 (High-wage and low-wage firms). *In a partial-segregation equilibrium, Mixed firms pay higher entry-level wages than Red firms: $w_{lm}^* > w_{lr}^*$.*

Since Mixed firms pay higher entry-level wages, we refer to them as *high-wage firms*. Blue workers strictly prefer to work in such firms, because they have a higher probability of promotion than Red workers. Red firms offer their young workers safe career paths, which can be interpreted as bureaucratic promotion rules, such as promotion by seniority, which are common in some government and nonprofit organizations. These firms are smaller than Mixed firms and pay low entry-level wages, so we refer to them as *low-wage firms*. While Red firms do not discriminate in hiring, they still end up with a fully Red workforce because only Red workers find it optimal to apply to them.

Subtle bias in promotions leads to inequality between Blue and Red workers, as the next corollary (of Proposition 4) shows.

Corollary 1 (Blue-Red inequality). *In equilibrium:*

- i. *Red workers are (weakly) less likely than Blue workers to work for firms offering risky career path contracts.*
- ii. *In Mixed firms, Blue workers have higher promotion probabilities and average wages than Red workers.*
- iii. *Blue workers have higher utility than Red workers.*

Corollary 1 shows that the Blue-Red inequality could be driven by both sorting and unequal treatment within firms. These two channels are linked. Because Red workers anticipate lower promotion probabilities than Blue workers in large, high-wage firms, Red workers are less likely to apply to these firms. Such results are consistent with the evidence in Card, Cardoso, and Kline (2016), who study gender gaps in employment and earnings across firms and find evidence of “(...) a sorting channel that arises if women are less likely to be employed at higher-wage firms, and a bargaining channel that arises if women obtain a smaller share of the surplus associated with their job.” Similarly, when investigating the recent widening of the gender earnings gap, Goldin et al. (2017) show that this widening “is split between men’s greater ability or preference to move to higher paying firms and positions and their better facility to advance within firms” (p. 114). Our model jointly explains both these facts and the emergence of high- and low-wage firms.

in Condition 3 and select the partial segregation equilibrium. Region II.B in Figure 1 reflects the parameter values for which full segregation is the only equilibrium.

3.6 Effect of Bias on Equilibrium Outcomes

To study the effect of the bias on the equilibrium outcomes, we focus here on the equilibrium with ex-post heterogeneous firms, that is, the case of $\beta \geq \beta_{max}$.²⁰ Since the equilibrium is characterized by a system of nonlinear equations, closed-form comparative statics typically cannot be obtained. We thus illustrate the equilibrium relationships using the same parametrization as in Figure 1.

Proposition 5 shows that firm heterogeneity is related to the size of the bias. Thus, we first focus on the effect of the bias on the differences between Red and Mixed firms. Figure 2 presents equilibrium wages in Red and Mixed firms, as well as the size (i.e., span of control) of Mixed firms, n_m^* , as functions of β . All panels show two equilibrium regions: partial segregation and full segregation (Regions II.A and II.B in Figure 1). The equilibrium values do not vary with β in the full-segregation region. This makes sense; under full segregation, the bias does not affect the firms' marginal conditions. In that region, the bias matters only through its off-equilibrium effect: Red workers do not apply to Blue-only firms because they expect promotions to be biased against them. It is this belief that allows Red firms to exist; these small firms absorb the supply of cheap Red workers.

Panel A shows wages in Red firms. Under partial segregation, as the bias increases, the outside option of Red workers, U_r^* , deteriorates, thereby lowering their bargaining power. Consequently, both entry-level and top wages decline as β increases.

Proposition 6 shows that Mixed firms pay higher entry-level wages than Red firms. Panel B shows that the same is true for top-level wages: $w_{hm}^* > w_{hr}^*$. Thus, Mixed firms are typically high-wage firms at all hierarchical levels. Panel B also shows that in Mixed firms, both the entry-level wage and the top wage *increase* with the bias (under partial segregation), which may seem counterintuitive. As firms become more biased against Red workers, the equilibrium utility U_r^* falls, which, from (20), implies that entry-level wages increase. At the same time, as workers become less powerful, the firms respond by expanding hiring to extract surplus from more workers. Panel C of Figure 2 indeed shows that Mixed firm size, or the *span of control*, n_m^* , increases with the bias, lowering the probability of promotion. To compensate for the lower probability of promotion, Mixed firms increase wages. This effect no longer holds under full segregation, as Mixed firms no longer attract Red workers.

²⁰If $\beta < \beta_{min}$, we have a corner solution where $s^* = 1$ (i.e., all firms are Mixed). In terms of predictions, this equilibrium is quite similar to the benchmark case with $\beta = 0$, although it also creates substantial inequality between Blue and Red agents. For completeness, we offer a full analysis of this equilibrium in the Internet Appendix.




Figure 2: Wages and Employment in Partial and Full-segregation Equilibrium.

Note: Panel A presents equilibrium wages in Red firms for the top job (w_{hr}^* , solid Red lines) and entry-level job (w_{lr}^* , dashed Red lines). Panel B shows corresponding wages in Mixed firms (w_{hm}^* , solid Blue lines; w_{lm}^* , dashed Blue lines). Panel C presents the equilibrium span of control in Mixed firms (n_m^*). Panel D shows average per-firm wages in Mixed firms ($\bar{w}_m$, solid Blue lines) and Red firms ($\bar{w}_r$, dashed Red lines). All panels plot outcomes as functions of subtle bias β , with $\theta = 10.0$, $R_l = 0.75$, $\alpha_b = 0.5$, and $e = 2.0$.

Finally, Panel D of Figure 2 presents the average per-firm wages in Mixed and Red firms, computed as

$$\bar{w}_m = \frac{n_m^* w_{lm}^* + w_{hm}^*}{n_m^* + 1} \text{ and } \bar{w}_r = \frac{w_{lr}^* + w_{hr}^*}{2}. \quad (23)$$

We see that, also on average, Mixed firms are “high-wage firms” and Red firms are “low-wage firms.”




Figure 3: Share of Blue workers in Mixed firms and share of Mixed firms in the economy.

Note: Panel A presents the equilibrium share of Blue workers in Mixed firms, q^* . Panel B shows the equilibrium share of Mixed firms in the economy, s^* . All panels plot outcomes as functions of subtle bias β , with $\theta = 10.0$, $R_l = 0.75$, $\alpha_b = 0.5$, and $e = 2.0$.

Figure 3 presents the worker composition of Mixed firms (Panel A) and the firm type composition of the economy (Panel B) as a function of subtle bias. Panel A shows that as the bias increases, the high-wage Mixed firms become increasingly dominated by Blue workers. Panel B shows that the share of Mixed firms in the economy decreases with the bias. Thus, as bias rises, high-wage firms become less common but (inefficiently) larger, hiring many favored agents and a minority of unfavored agents. By contrast, the fraction of smaller, bureaucratic, low-wage firms increases in the economy.

Figure 4 shows the effect of β on workers' utilities. Unsurprisingly, a higher bias against Red workers decreases their equilibrium utility.²¹ Because Mixed firms choose contracts that bind Red workers' participation constraints under partial segregation, weaker outside options for Red workers lead to less favorable terms for all workers. Although Blue workers are favored by the bias, as β increases, the fraction of Blue in Mixed firms, q^* , also increases, as well as firm size, n_m^* . The latter effect outweighs the positive bias effect, implying that Blue workers' utility also decreases as β increases. For sufficiently high β , both Blue and Red workers are worse off than in the benchmark equilibrium (shown as horizontal lines in Figure 4).

²¹We formally show this result as part of the proof of Proposition 5.




Figure 4: Utilities of Blue and Red workers, U_b^* and U_r^* .

Note: This figure presents the equilibrium utilities of Blue workers (solid blue lines), U_b^* and Red workers (dashed red lines), U_r^* , as functions of subtle bias, β , for $\theta = 10.0$, $R_l = 0.75$, $\alpha_b = 0.5$, and $e = 2.0$. The horizontal dotted line represents the equilibrium utility, $U_b^* = U_r^* = U^*$, in the benchmark case, $\beta = 0$.

The following proposition shows the effect of the bias on profits.

Proposition 7 (Profit increases with bias). *For $\beta \geq \beta_{max}$, equilibrium profits are weakly increasing in β .*

In an equilibrium with ex-post heterogeneous firms, Red and Mixed contracts must yield the same profits. As the bias increases, Red workers' bargaining power decreases, and profits increase with β .

Figure 5 shows that not only profits increase with the bias, but that firms may prefer the (inefficient) “biased equilibrium” to the (efficient) benchmark equilibrium.²² Indeed, when β is sufficiently high, equilibrium profits exceed those in the benchmark case, as indicated by the horizontal line in Figure 5. Intuitively, the subtle bias acts as an implicit collusion device in the labor market: by depressing the outside options of Red workers, it allows firms to capture a larger share of a shrinking surplus. Perhaps unexpectedly, the model can rationalize postmodern theories that view discrimination against minorities as an implicit capitalist plot to increase profits by extracting more “surplus value” from workers. The surprising part is that we obtain this result in a model with perfect competition and rational expectations.

²²In Figure 5, we set $R_h = 4.0$. Parameter R_h does not affect equilibrium outcomes and thus can assume any arbitrary positive level to achieve strictly positive profits.




Figure 5: Firm profit, π^* .

Note: This figure presents the equilibrium profit of Mixed and Red firms, π^* , as a function of subtle bias, β , for $\theta = 10.0$, $R_l = 0.75$, $R_h = 4.0$, $\alpha_b = 0.5$, and $e = 2.0$. The horizontal dotted line represents the equilibrium profit of Mixed firms, π^* , in the benchmark case, $\beta = 0$.

3.7 Endogenous Firm Entry

Until now, we have kept the number of firms, F , constant. In price-theory parlance, our equilibrium results can be interpreted as “short run,” that is, before firms make entry and exit decisions. We now assume that firms may enter the sector by paying a fixed operating cost in each period, $\kappa(F) = a + bF$, where $a, b \geq 0$, and F is the number of firms in that period. This specification captures several possibilities. If $a = b = 0$, there are no fixed operating costs. If $a > 0$ and $b = 0$, operating costs are independent of the number of firms in the sector. Finally, if $b > 0$, such costs increase with the number of firms. There are several possible interpretations in this case. First, if firms are heterogeneous in their operating costs, we can interpret $\kappa(F)$ as the operating cost of the “marginal firm,” that is, the least productive firm that enters the sector. Alternatively, keeping the assumption that all firms are ex ante identical, $\kappa(F)$ may increase with F because of increased competition in the product market (which we currently model only in reduced form), or because of increased general expenses (such as R&D, human resources, and marketing) that increase with the intensity of competition in both labor and product markets.

Firms enter only if the expected profit net of operating costs is non-negative. Thus, the equilibrium number of firms is determined by:

$$a + bF^* = \pi^*(F^*), \quad (24)$$

where $\pi^*(F^*)$ is the equilibrium (gross) profit when firm owners anticipate that F^* firms enter. Figure 6 illustrates how the equilibrium F^* is determined. The equilibrium gross profit is a piecewise-

defined, regime-dependent function of the number of firms, F (solid blue), while the operating cost function is upward-sloping (dashed black line). Their intersection determines the equilibrium number of firms, F^* , and the corresponding equilibrium profit, π^* .




Figure 6: Equilibrium Number of Firms, F^* , and Profit, π^* .

Note: This figure presents the equilibrium gross profit, $\pi^*(F)$, as a function of F for $\theta = 10$, $\beta = 0.7$, $R_l = 0.7$, $R_h = 4$, $\alpha_b = 0.5$, $E = 200$. It also shows the operating cost function (black dashed line), for $a = 3.2$ and $b = 0.0022$. Equilibrium values (F^* , π^*) occur when the two functions intersect.

Figure 6 illustrates that gross profit is decreasing in F in Region I and increasing in Region II.A (partial segregation). In contrast, in Region II.B, gross profit is locally flat: additional entry takes the form of firms offering safe-career-path contracts that attract only Red workers from unemployment, whose outside option remains fixed at zero. As a result, entry does not bid up wages or otherwise affect profits, leaving $\pi^*(F)$ unchanged over this range of F . Importantly, as the figure illustrates, for a stable equilibrium in partial segregation to exist, we must have $b > 0$.²³

As it can be inferred from Figure 6, an increase in operating costs—due to an increase in either a or b —reduces the number of firms operating in the sector. Thus, lower operating costs favor the equilibrium with heterogeneous firms (Region II). But the entry of new firms does not offset the inefficiencies caused by the bias. Entry in Region II.A implies *fewer* Mixed firms, which become even larger and hire more Blue workers. Thus, as q^* increases, the *effective bias* ($q^*\beta$) increases with firm entry.

Firm entry has surprising welfare effects. As Figure 7 illustrates, in Region I, workers' utilities increase with firm entry, while in Region II.A, workers' utilities *decrease* with entry. This result follows because, in an equilibrium with two firm types, the marginal firm entering the sector is a

²³The operating cost function may cross the gross profit function at most three times. For a stable equilibrium, the operating cost function must cross the profit function from below, which implies that at most two equilibria are stable. Our equilibrium selection Condition 3 implies that if there are two stable equilibria, the equilibrium in Region I is chosen.

Red firm, which is inefficiently small. Thus, the entry of organizationally inefficient firms shifts Red workers from Mixed firms to Red firms, destroying total surplus while reducing Red workers' bargaining power.




Figure 7: Utilities of Blue and Red workers, U_b^* and U_r^* .

Note: This figure presents the equilibrium utilities of Blue and Red workers, U_b^* and U_r^* , as functions of the number of firms, F , for $\theta = 10.0$, $\beta = 0.7$, $R_l = 0.7$, $R_h = 4.0$, $\alpha_b = 0.5$, $E = 200.0$.

Finally, note that in an equilibrium in Region II.A, an increase in β will increase the gross profit (see Proposition 7), thus attracting more firms to the sector. But, again, more firms do not attenuate the negative effects of the bias. Competition does not drive discriminating firms out of business in our model because biases are subtle, that is, they are not observable. In the Internet Appendix, we show that subtle bias persists in equilibrium even when firms can acquire a costly bias-reducing technology.

4 Conclusion

This paper presents a model of biased promotions where firm size, wages, and internal labor markets are endogenously determined in a competitive labor market equilibrium. By providing a tractable framework for the equilibrium consequences of biased promotions, the paper contributes to the broader literatures on employer discrimination, firm heterogeneity, and inequality.

Claudia Goldin's (2014) description of women being "*barred in subtle and more obvious ways from many high-prestige and high-income occupations*" maps into the two main ingredients in our model. The prestige and high-income of top positions are what workers compete for; the subtlety of the bias is what makes discrimination persistent. Without the complementarity between pay and prestige, internal labor markets would not exist. Without subtlety, entry by unbiased firms would

eliminate discrimination. Together, the two ingredients deliver partial segregation in equilibrium: Red workers concentrate in smaller, lower-paying Red firms that offer certain promotion, while Mixed firms grow larger, pay more, and reserve the bulk of their promotions for Blue workers. Goldin’s “female-dominated occupations”—teacher, nurse, librarian, social worker—have something of the flavor of our Red firms: orderly internal hierarchies, predictable advancement, and lower pay.

Several of our results go against the intuition that competitive labor markets discipline discrimination. Bias can raise firm profits because it depresses Red workers’ outside options, so the bias operates as an implicit coordination device on the labor demand side. Firm entry can amplify rather than dissipate the harm: when marginal entrants are inefficiently small Red firms, expected earnings fall for everyone. Perfect competition and rational expectations are sufficient to generate these outcomes; no belief distortions or coordination failures are required.

The model also speaks to a nascent empirical literature on amenities and gender earnings gaps (Morchio and Moser (2026); Caldwell, Haeghele, and Heining (2026)). When prestige is the scarce amenity, biased access to it produces wage and sorting patterns that resemble differential preferences over jobs. Our model illustrates that separating taste-based sorting from biased access to prestigious positions is a challenging empirical task.

A Appendix

Proof of Proposition 1. Probabilities adding up to one imply $p_r(n, 0) = \frac{1}{n}$. Because Axiom (i) implies $p_r(n, q)$ must be decreasing in q , without loss of generality, we can write $p_r(n, q) = \frac{1 - \omega_r(n, q)}{n}$, where $\omega_r(n, q) \in [0, 1]$ is increasing in q and $\omega_r(n, 0) = 0$. Similarly, because $p_b(n, 1) = \frac{1}{n}$, we can write $p_b(n, q) = \frac{1 + \omega_b(n, q)}{n}$, where $\omega_b(n, q) \in [0, n - 1]$ is decreasing in q and $\omega_b(n, 1) = 0$. Axiom (ii) then implies that $\omega_i(n, q) = \omega_i(q)$ for $i \in \{b, r\}$. Axiom (iii) implies that $\omega_i(q)$ is linear. We then write $\omega_r(q) = q\omega_r$ (there is no intercept because $\omega_r(0) = 0$) and $\omega_b(q) = (1 - q)\omega_b$ (because $\omega_b(1) = 0$). Probabilities must add up to 1: $q(1 + (1 - q)\omega_b) + (1 - q)(1 - q\omega_r) = 1$, which implies $\omega_b = \omega_r := \beta$. Finally, $p_r(2, \frac{1}{2}) = \frac{1 - \beta}{2} = \frac{1 - \frac{1}{2}\beta}{2}$, so $\beta = 2\beta_b$. $\square$

Proof of Proposition 2. We assume that at least one tight-labor-market equilibrium exists. Consider one such equilibrium and let U^* denote a worker’s lifetime utility in equilibrium. Profit maximization implies that firms solve the constrained maximization problem:

$$\max_{n, w_l, w_h, \lambda, \gamma} R_h - w_h + n(R_l - w_l) + \lambda \left(u(w_l) + \frac{1}{n} \theta u(w_h) - U^* \right) + \gamma(n - 1). \quad (\text{A.1})$$

Because $E > F$, in any equilibrium, firms must offer $n^* > 1$, thus we set $\gamma^* = 0$. With $\gamma^* = 0$, the problem is globally concave (see Ferreira and Nikolowa (2024)), and thus the first-order conditions

are sufficient to determine a unique maximum given U^* . Thus, all firms must offer the same contract. The first-order conditions are

$$\frac{\partial L}{\partial w_l} = -n^* + \lambda^* u'(w_l^*) = 0 \quad (\text{A.2})$$

$$\frac{\partial L}{\partial w_h} = -1 + \lambda^* \frac{1}{n^*} \theta u'(w_h^*) = 0 \quad (\text{A.3})$$

$$\frac{\partial L}{\partial n} = R_l - w_l^* - \lambda^* \frac{1}{(n^*)^2} \theta u(w_h^*) = 0 \quad (\text{A.4})$$

$$\frac{\partial L}{\partial \lambda} = u(w_l^*) + \frac{1}{n^*} \theta u(w_h^*) - U^* = 0 \quad (\text{A.5})$$

Conditions (A.2) and (A.3) imply that the marginal utilities must be equal:

$$u'(w_l^*) = \theta u'(w_h^*). \quad (\text{A.6})$$

Use (A.4) to isolate λ^* and replace it in (A.2) to find

$$R_l - w_l^* = \frac{\theta u(w_h^*)}{n^* u'(w_l^*)}. \quad (\text{A.7})$$

Using (A.5) we can rewrite (A.7) as

$$R_l - w_l^* = \frac{U^* - u(w_l^*)}{u'(w_l^*)}. \quad (\text{A.8})$$

Because the problem is globally concave, there is a unique solution for (A.8), $w_l(U^*)$. From (A.8) we derive $\frac{\partial w_l}{\partial U^*} = \frac{u'(w_l)}{u''(w_l)(U^* - u(w_l))}$, it therefore follows that $\frac{\partial w_l}{\partial U^*} < 0$.

After finding w_l^* , we find w_h^* from (A.6), and then n^* from (A.5):

$$n^* = \frac{\theta u(w_h^*)}{U^* - u(w_l^*)} = \frac{E}{F}. \quad (\text{A.9})$$

We find the equilibrium by solving for U^* in

$$u(w_l(U^*)) + \frac{F}{E} \theta u(w_h(U^*)) = U^*. \quad (\text{A.10})$$

Because $w_l'(U^*) < 0$ and $w_h'(U^*) < 0$, there is only one solution to (A.10). $\square$

Proof of Proposition 3. The argument in the text implies that no firm can have a mixed workforce in equilibrium. Thus, firms must be fully segregated: “Blue firms” and “Red firms.” Both types must have the same equilibrium profit. Suppose a worker of type i has a higher equilibrium utility than type j . Since overt discrimination is permitted, a firm currently hiring type- i workers can deviate and offer a similar contract with slightly lower wages to type- j workers, who would accept. We conclude that both worker types must have the same utility in equilibrium. Thus, all firms face the same participation constraint. Proposition 2 then implies that both types of firms must offer identical contracts (except for the exclusion of one worker type). Because of full segregation, all

workers in a firm face the same probability of promotion, which is $1/n^*$. Thus, they must be offered the unique contract described in Proposition 2. $\square$

Proof of Lemma 1. Suppose C^* is a tight-labor market equilibrium with full segregation and homogeneous firms. If $U_r^*(C^*) \neq U_b^*(C^*)$, firms that employ workers with the highest utility have lower profits, which is not optimal for them since they could target the low-utility workers instead. If $U_r^*(C^*) = U_b^*(C^*)$, all firms face the same maximization problem. Thus, the optimal number of workers per firm must be unique, n^* . $n^* = 1$ cannot be a tight-labor-market equilibrium because $E > F$. Therefore, all firms must offer $n^* > 1$. In this case, if firms are segregated, Blue workers would prefer to direct their search to firms where $q(c_\tau^*, C_{-\tau}^*) = 0$, because $p_b(n, q)$ is decreasing in q . Thus, these firms must select some Blue workers, which is a contradiction. $\square$

Proof of Proposition 4. Consider first a fully tight labor market equilibrium. Suppose that $\exists \tau \in [0, 1]$ such that $q_\tau^* := q(c_\tau^*, C_{-\tau}^*) \in (0, 1)$. That is, a strictly Mixed firm exists.²⁴ Suppose this firm is such that $n_\tau^* = 1$ (i.e., a safe career path). Then, both types of workers have the same utility under c_τ^* . Because $E > F$, under a fully tight labor market, there must be at least another firm such that $n_{\tau'}^* > 1$. If firm τ' is either strictly Mixed or all Red, Red workers must be indifferent between $c_{\tau'}^*$ than c_τ^* . But then Blue workers would be strictly better off under $c_{\tau'}^*$ than c_τ^* , implying that these cannot both be equilibrium contracts. If firm τ' is instead all Blue, we must have that the Blue worker is indifferent between c_τ^* and $c_{\tau'}^*$, and also that $\pi(c_\tau^*) = \pi(c_{\tau'}^*)$, otherwise the firm with the lowest profit could deviate and mimic the one with the highest profit. But then both firms face the same maximization problem as in (A.1), of which $n_{\tau'}^* > n_\tau^* = 1$ cannot both be optimal solutions (given strict global concavity). We conclude that the strictly Mixed firm must have $n_\tau^* > 1$.

We now show that all strictly Mixed firms offer the same contract, in which Red's participation constraint binds while Blue's is slack. Consider a strictly Mixed firm such that $n_\tau^* > 1$. Conditions 1-2 and profit maximization imply

$$\pi(c_\tau^*) = \max_{w_{l\tau}, w_{h\tau}, n_\tau} R_h - w_{h\tau} + n_\tau(R_l - w_{l\tau}) \quad (\text{A.11})$$

subject to

$$u(w_{l\tau}) + \frac{1 - q_\tau^* \beta}{n_\tau} \theta u(w_{h\tau}) \geq \bar{U}_r(c_\tau^*, C_{-\tau}^*) \quad (\text{A.12})$$

$$u(w_{l\tau}) + \frac{1 + (1 - q_\tau^*) \beta}{n_\tau} \theta u(w_{h\tau}) \geq \bar{U}_b(c_\tau^*, C_{-\tau}^*) \quad (\text{A.13})$$

$$n_\tau \geq 1. \quad (\text{A.14})$$

²⁴ As discussed in Footnote 10, we assume that the equilibrium is such that $(1 - q_\tau^*) \beta \leq n_\tau^* - 1$. If $(1 - q_\tau^*) \beta > n_\tau^* - 1$, the equilibrium probabilities of promotion must be instead $p_b(n_\tau^*, q_\tau^*) = 1$ and $p_r(n_\tau^*, q_\tau^*) = \frac{1 - n_\tau^* q_\tau^*}{n_\tau^* - n_\tau^* q_\tau^*}$. The analysis of this case is similar and thus omitted for brevity.

One of (A.12) and (A.13) must bind; otherwise, firm τ can always increase n to increase profit.

Suppose the Blue constraint (A.13) binds. Consider first a deviation contract c_τ^d that maximizes profit while ignoring the Blue constraint. Let $q_\tau^d := q(c_\tau^d, C_{-\tau}^*)$ denote the belief after this deviation. Because this is an “optimal deviation,” the Red constraint must bind (otherwise the firm can increase n_τ^d , which relaxes (A.14) and increases profit), implying $U_r(c_\tau^d; q_\tau^d) = \bar{U}_r(c_\tau^*, C_{-\tau}^*)$, where $U_i(c; q)$ denotes the utility of a type- i agent that accepts contract c under belief q . Note that in equilibrium, the Blue worker has strictly higher utility than the Red worker:

$$\bar{U}_b(c_\tau^*, C_{-\tau}^*) = U_b(c_\tau^*; q_\tau^*) = U_r(c_\tau^*; q_\tau^*) + \frac{\beta}{n_\tau^*} \theta u(w_{h\tau}^*) > \bar{U}_r(c_\tau^*, C_{-\tau}^*). \quad (\text{A.15})$$

If c_τ^d does not violate the Blue constraint, Condition 2 implies that $q_\tau^d = q_\tau^*$, and thus profit maximization implies that $\pi(c_\tau^d) \leq \pi(c_\tau^*) \Rightarrow$ the deviation is not profitable. Thus, c_τ^d must violate the Blue constraint. Condition 2 now implies that $q_\tau^d = 0$, which means only Red workers are matched to the deviating firm. There are two cases to consider.

Case 1. Suppose that (A.14) also binds (in addition to the Red constraint), i.e., $n_\tau^d = 1$. Because $U_r(c_\tau^d; q_\tau^d) = \bar{U}_r(c_\tau^*, C_{-\tau}^*)$, Blue workers strictly prefer some contract in $C_{-\tau}^*$ to c_τ^d , which confirms that no Blue worker applies to contract c_τ^d . Thus, c_τ^d is a profitable deviation for the firm $\Rightarrow$ the Blue constraint (A.13) cannot bind in an equilibrium where firm τ hires both Red and Blue workers.

Case 2. Suppose instead that, under the optimal deviation, $n_\tau^d > 1$. We then have

$$U_b(c_\tau^d; q_\tau^d) = U_r(c_\tau^d; q_\tau^d) + \beta \theta \frac{u(w_{h\tau}^d)}{n_\tau^d} = \bar{U}_r(c_\tau^*, C_{-\tau}^*) + \beta \frac{R_l - w_{l\tau}^d}{(1 - q_\tau^d \beta)} u'(w_{l\tau}^d), \quad (\text{A.16})$$

where the last step follows from the first-order conditions when the Red constraint binds. When the Red constraint binds, $w_{l\tau}^d$ is independent of q_τ^d . Thus, $U_b(c_\tau^d; q)$ increases with q , which implies $U_b(c_\tau^d; 0) < \bar{U}_b(c_\tau^*, C_{-\tau}^*)$, which confirms that no Blue worker applies to contract c_τ^d . Thus, c_τ^d is a profitable deviation for the firm $\Rightarrow$ the Blue constraint (A.13) cannot bind in an equilibrium where firm τ hires both Red and Blue workers.

Cases 1 and 2 jointly imply that for a strictly Mixed firm, the Blue constraint must not bind, and thus the Red constraint must bind. Since a strictly Mixed firm must have more than a unit mass of workers, there is only one contract that maximizes profit for a given q_τ^* . Suppose there is another risky-career-path firm with $q_{\tau'}^* \in [0, 1]$ such that $q_{\tau'}^* \neq q_\tau^*$. At least one type of worker must be indifferent between c_τ^* and $c_{\tau'}^*$. But then the firm with the highest q must pay higher wages and/or hire fewer workers, implying $\pi(c_\tau^*) \neq \pi(c_{\tau'}^*) \Rightarrow$ this cannot be an equilibrium. Thus, there is only one equilibrium pair (c_τ^*, q_τ^*) for strictly Mixed firms, implying that all such firms are identical, and there are no risky-career-path firms with only Blue or only Red workers. We denote the optimal contract for a strictly Mixed firm by $c_m^* = (n_m^*, w_{lm}^*, w_{hm}^*)$. Note that $U_b(c_m^*, C_{-m}^*) > U_r(c_m^*, C_{-m}^*)$.

Suppose contract c_m^* coexists with a safe-career-path contract $c_r^* = (1, w_{lr}^*, w_{hr}^*)$. Equilibrium

implies $\pi(c_m^*) = \pi(c_r^*)$, and worker types who apply to both contracts must be indifferent between them. A safe-career-path contract offers the same utility to Blue or Red workers, since promotion is guaranteed. Because $U_b(c_m^*, C_{-m}^*) > U_r(c_m^*, C_{-m}^*)$, c_r^* can only attract Red workers, who must be indifferent between c_m^* and c_r^* . Trivially, no other safe-career-path contract can exist, as it would have a different profit than c_r^* . We refer to the firms offering c_r^* as Red firms, because only Red workers apply to such a contract. By continuity, the same argument applies to the case where $q_\tau^* = 1$.

If $q_\tau^* = 0$, Lemma 1 implies that all Blue workers must choose a safe-career-path contract. Red workers must weakly prefer c_τ^* to the safe career path, and thus have equilibrium utility no lower than Blue workers, since both worker types enjoy the same utility under the same safe career path contract. But this is not possible because a Blue worker who deviates and chooses c_τ^* must obtain higher utility than a Red worker (because of the promotion bias).

Finally, under a fully tight labor market, a risky career path contract must exist, otherwise not all workers will be employed.

We then conclude that under a fully tight labor market: (i) a Mixed-firm contract c_τ^* exists and is unique within its class; (ii) Blue workers must apply only to c_τ^* ; and (iii) if another contract exists, it must be a safe career path (c_r^*), attracting only Red workers.

We now consider the case of a partially tight labor market equilibrium. If $U_r^* > U_b^* = 0$, this implies that the labor market for Red must be tight. However, for any contract with $n_\tau^* \geq 1$, the utility of a Blue worker is higher than or equal to the Utility of a Red worker; therefore, we cannot have $U_r^* > U_b^* = 0$ in equilibrium.

Thus, a partially tight equilibrium requires $U_b^* \geq U_r^* = 0$, with the market tight for Blue and slack for Red. Consider a contract c_m^* with $q_m^* < 1$, as defined above, where the Red's constraint binds. Let c_r^* denote the optimal safe career path attracting only Red workers when $U_r^* = 0$. If $\pi(c_m^*) \geq \pi(c_r^*)$, we are back to the case discussed above, where contracts (c_m^*, c_r^*) are the only ones admissible in equilibrium, and the mass of firms offering c_r^* is nonzero if and only if $\pi(c_m^*) = \pi(c_r^*)$. If instead $\pi(c_m^*) < \pi(c_r^*)$, then no firm would offer c_m^* . However, if all firms offer a safe career contract, then the market cannot be partially tight. Thus, the only remaining possibility is to have contract c_r^* attracting only Red workers, and a contract c_m^* with $q_m^* = 1$ in which the participation constraint of Blue workers binds. In this case, equilibrium requires $U_r(c_m^*, c_r^*) < 0 = U_r(c_r^*, c_m^*)$ (confirming that $q_m^* = 1$), and $\pi(c_m^*) = \pi(c_r^*)$. $\square$

Proof of Proposition 5. From Proposition 4, we know that, for a given β , some firms must offer a Mixed contract $c_m^*(\beta)$ in equilibrium. For given β and q , define $\Delta(\beta, q) \equiv \pi(c_m^*(\beta, q)) - \pi(c_r^d(\beta, q))$, where $c_m^*(\beta, q)$ is a “conjectured” equilibrium contract for a Mixed firm under belief q , and $c_r^d(\beta, q)$ is the optimal deviation contract a firm would offer if unilaterally deviating from this equilibrium. As discussed in the proof of Proposition 4, this optimal deviation must be a

safe career path contract (which may or may not be $c_r^*(\beta, q)$).

First, we show that for sufficiently low β , firms are homogeneous. At the limit, i.e., for $\beta = 0$, the maximization problem is:

$$\max_{w_{lm}, w_{hm}, n_m} R_h - w_{hm} + n_m(R_l - w_{lm}) \quad (\text{A.17})$$

subject to

$$\begin{cases} u(w_{lm}) + \frac{1}{n_m} \theta u(w_{hm}) \geq U_r^* \\ n_m \geq 1 \end{cases} \quad (\text{A.18})$$

For a given U_r^* , the resulting optimal contract is unique (see Proposition 2), and in order to have a tight labor market equilibrium, the parameters are such that $n_m^* > 1$, implying $q^* = \alpha_b$. It then follows that the profit from offering a risky career path is higher than the profit from offering a safe career path, and therefore $\Delta(0, \alpha_b) = \pi(c_m^*(0, \alpha_b)) - \pi(c_r^d(0, \alpha_b)) > 0$. The Maximum Theorem implies that $\Delta(\beta, \alpha_b)$ is continuous in β , and thus there exists $\beta_{min} \equiv \inf\{\beta : \Delta(\beta, \alpha_b) \leq 0\} \in (0, 1]$, implying that $\Delta(\beta, \alpha_b) > 0$ for all $\beta < \beta_{min}$, in which case all firms offer $c_m^*(\beta, \alpha_b)$ in equilibrium. This proves Part 1.

Define $\beta_{max} \equiv \sup\{\beta : \Delta(\beta, \alpha_b) > 0\} \in (0, 1]$. It then follows that for $\beta > \beta_{max}$, the case where all firms are Mixed is not an equilibrium. This proves Part 2. $\square$

Proof of Proposition 6: The equilibrium contract in Red firms must satisfy the marginal condition $u'(w_{lr}^*) = \theta u'(w_{hr}^*)$, and the partial segregation equilibrium contract in Mixed firms must be such that $u'(w_{lm}^*) = \theta(1 - q^*\beta)u'(w_{hm}^*)$. We rearrange the marginal conditions:

$$\frac{u'(w_{lr}^*)}{u'(w_{lm}^*)} = \frac{u'(w_{hr}^*)}{(1 - q^*\beta)u'(w_{hm}^*)}.$$

Suppose $w_{lr}^* \geq w_{lm}^*$. Then, we have

$$\frac{u'(w_{lr}^*)}{u'(w_{lm}^*)} = \frac{u'(w_{hr}^*)}{(1 - q^*\beta)u'(w_{hm}^*)} \leq 1 \Rightarrow u'(w_{hr}^*) < u'(w_{hm}^*) \Rightarrow w_{hr}^* > w_{hm}^*.$$

Because in equilibrium, a Red worker must be indifferent between the two contracts,

$$u(w_{lm}^*) + p_r(n_m^*, q^*)\theta u(w_{hm}^*) = u(w_{lr}^*) + \theta u(w_{hr}^*)$$

or, rearranging,

$$u(w_{lm}^*) - u(w_{lr}^*) = \theta[u(w_{hr}^*) - p_r(n_m^*, q^*)u(w_{hm}^*)], \quad (\text{A.19})$$

where $p_r(n_m^*, q^*) = \frac{1 - q^*\beta}{n_m^*} < 1$. If $w_{lr}^* \geq w_{lm}^*$ and $w_{hr}^* > w_{hm}^*$, the left-hand side of (A.19) is non-positive and the right-hand side is strictly positive, which is a contradiction. Thus, we must have $w_{lr}^* < w_{lm}^*$. $\square$

Proof of Corollary 1: In any (partially or fully) tight labor market equilibrium, Blue workers work

in firms offering risky career paths with probability 1. Red workers work in firms offering risky career paths with probability

$$\frac{s^* n_m^* (1 - q^*)}{s^* n_m^* (1 - q^*) + (1 - s^*)} \leq 1.$$

In Mixed firms, Blue workers' probability of promotion is

$$p_b(n_m^*, q^*) = \frac{1 + (1 - q^*)\beta}{n_m^*} > \frac{1 - q^*\beta}{n_m^*} = p_r(n_m^*, q^*).$$

Expected wages are $w_{lm}^* + p_i(n_m^*, q^*)w_{hm}^*$ and therefore it immediately follows that the expected wages of Blue workers are higher. Finally, from the proof of Proposition 4, $U_b(c_m^*, C_{-m}^*) = U_r(c_m^*, C_{-m}^*) + \frac{\beta}{n_m^*} \theta u(w_{hm}^*)$, that is $U_b(c_m^*, C_{-m}^*) > U_r(c_m^*, C_{-m}^*)$, for any $\beta > 0$. $\square$

Proof of Proposition 7: For $\beta \geq \beta_{max}$, $\pi(c_r^*) = \pi(c_m^*)$. The optimal c_r^* contract is such that:

$$u'(w_{lr}) = \theta u'(w_{hr}) \quad (\text{A.20})$$

$$u(w_{lr}) + \theta u(w_{hr}) = U_r^* \quad (\text{A.21})$$

From the participation constraint and condition (A.20):

$$\begin{aligned} u'(w_{lr}) \left(\frac{\partial w_{lr}}{\partial \beta} + \frac{\partial w_{hr}}{\partial \beta} \right) &= \frac{\partial U_r^*}{\partial \beta} \\ \frac{\partial w_{lr}}{\partial \beta} + \frac{\partial w_{hr}}{\partial \beta} &= \frac{\partial U_r^*}{\partial \beta} \frac{1}{u'(w_{lr})} \end{aligned} \quad (\text{A.22})$$

To show that $\frac{\partial U_r^*}{\partial \beta} \leq 0$, rewrite and simplify the profit equalization condition as

$$LC_r \equiv w_{lr} + w_{hr} - R_l = w_{hm} - \frac{u(w_{hm})}{u'(w_{hm})} \equiv LC_m. \quad (\text{A.23})$$

Differentiate (A.23) with respect to β :

$$\begin{aligned} \frac{\partial LC_r}{\partial U_r} \frac{\partial U_r}{\partial \beta} &= \frac{\partial LC_m}{\partial U_r} \frac{\partial U_r}{\partial \beta} + \frac{\partial LC_m}{\partial \beta} + \frac{\partial LC_m}{\partial q} \frac{\partial q}{\partial \beta} \\ \Leftrightarrow \left(\frac{\partial LC_r}{\partial U_r} - \frac{\partial LC_m}{\partial U_r} \right) \frac{\partial U_r}{\partial \beta} &= \frac{\partial LC_m}{\partial \beta} + \frac{\partial LC_m}{\partial q} \frac{\partial q}{\partial \beta} \end{aligned} \quad (\text{A.24})$$

Use Red workers' participation constraint and condition $\theta u'(w_{hr}) = u'(w_{lr})$ to get

$$u'(w_{lr}) \frac{\partial w_{lr}}{\partial U_r} + \theta u'(w_{hr}) \frac{\partial w_{hr}}{\partial U_r} = 1 \Rightarrow \frac{\partial w_{lr}}{\partial U_r} + \frac{\partial w_{hr}}{\partial U_r} = \frac{1}{u'(w_{lr})}. \quad (\text{A.25})$$

From (A.25) and the definition of LC_r it follows that:

$$\frac{\partial LC_r}{\partial U_r} = \frac{\partial w_{lr}}{\partial U_r} + \frac{\partial w_{hr}}{\partial U_r} = \frac{1}{u'(w_{lr})}. \quad (\text{A.26})$$

Differentiate LC_m with respect to U_r :

$$\begin{aligned}
\frac{\partial LC_m}{\partial U_r} &= \frac{u''(w_{hm})u(w_{hm})}{u'(w_{hm})^2} \frac{\partial w_{hm}}{\partial U_r} \\
&= \frac{u''(w_{hm})u(w_{hm})}{u'(w_{hm})^2} \frac{1}{\theta(1-q\beta)u''(w_{hm})} \frac{\theta(1-q\beta)u'(w_{hm})}{(U_r - u(w_{lm}))} \\
&= \frac{1}{u'(w_{hm})} \frac{u(w_{hm})}{(U_r - u(w_{lm}))} \frac{\theta(1-q\beta)}{\theta(1-q\beta)} \\
&= \frac{n_m}{u'(w_{lm})}.
\end{aligned} \tag{A.27}$$

From (A.26) and (A.27), we have:

$$\frac{\partial LC_r}{\partial U_r} - \frac{\partial LC_m}{\partial U_r} = \frac{1}{u'(w_{lr})} - \frac{n_m}{u'(w_{lm})} < 0. \tag{A.28}$$

The inequality in (A.28) follows from: (i) $w_{lm} > w_{lr}$ and $u''(\cdot) < 0$ imply $u'(w_{lm}) < u'(w_{lr})$, and (ii) $n_m > 1$.

β affects only the contract offered by Mixed firms. For given U_r , w_{lm} is unaffected by β . From $\theta(1-q\beta)u'(w_{hm}) = u'(w_{lm})$, we have:

$$\frac{\partial w_{hm}}{\partial \beta} = \frac{\theta q u'(w_{hm})}{\theta(1-q\beta)u''(w_{hm})} \text{ and } \frac{\partial w_{hm}}{\partial q} = \frac{\theta \beta u'(w_{hm})}{\theta(1-q\beta)u''(w_{hm})}. \tag{A.29}$$

$$\begin{aligned}
\frac{\partial LC_m}{\partial \beta} &= \frac{\partial LC_m}{\partial w_{hm}} \frac{\partial w_{hm}}{\partial \beta} \\
&= \frac{\theta q u(w_{hm})}{u'(w_{lm})} > 0
\end{aligned} \tag{A.30}$$

$$\begin{aligned}
\frac{\partial LC_m}{\partial q} &= \frac{\partial LC_m}{\partial w_{hm}} \frac{\partial w_{hm}}{\partial q} \\
&= \frac{\theta \beta u(w_{hm})}{u'(w_{lm})} > 0
\end{aligned} \tag{A.31}$$

Using the labor market clearing conditions $s^* n_m^* q^* = \alpha_b e$ and $s^* n_m^* + 1 - s^* = e$, we get:

$$\begin{aligned}
q &= \frac{\alpha_b e}{e-1} \left(1 - \frac{1}{n_m^*} \right) \\
&= \frac{\alpha_b e}{e-1} \left(1 - \frac{u'(w_{hm})(R_l - w_{lm})}{u(w_{hm})} \right).
\end{aligned} \tag{A.32}$$

For a given U_r , we have:

$$\begin{aligned}
\frac{\partial q}{\partial \beta} &= \frac{\alpha_b e}{e-1} \frac{\beta}{(1-q\beta)n_m} \left(-1 + \frac{u'(w_{hm})^2}{u''(w_{hm})u(w_{hm})} \right) \frac{\partial q}{\partial \beta} + \frac{\alpha_b e}{e-1} \frac{q}{(1-q\beta)n_m} \left(-1 + \frac{u'(w_{hm})^2}{u''(w_{hm})u(w_{hm})} \right) \\
&= \frac{\frac{\alpha_b e}{e-1} \frac{q}{(1-q\beta)n_m} \left(-1 + \frac{u'(w_{hm})^2}{u''(w_{hm})u(w_{hm})} \right)}{1 - \frac{\alpha_b e}{e-1} \frac{\beta}{(1-q\beta)n_m} \left(-1 + \frac{u'(w_{hm})^2}{u''(w_{hm})u(w_{hm})} \right)}
\end{aligned} \tag{A.33}$$

We can now rewrite (A.28) as follows:

$$\begin{aligned}
& \left(\frac{1}{u'(w_{lr})} - \frac{n_m}{u'(w_{lm})} \right) \frac{\partial U_r}{\partial \beta} = \frac{\theta u(w_{hm})}{u'(w_{lm})} \left(q + \beta \frac{\partial q}{\partial \beta} \right) \\
& \Leftrightarrow \left(\frac{1}{u'(w_{lr})} - \frac{n_m}{u'(w_{lm})} \right) \frac{\partial U_r}{\partial \beta} = \frac{\theta q u(w_{hm})}{u'(w_{lm})} \left(1 + \frac{\frac{\alpha_b e}{e-1} \frac{\beta}{(1-q\beta)n_m} \left(-1 + \frac{u'(w_{hm})^2}{u''(w_{hm})u(w_{hm})} \right)}{1 - \frac{\alpha_b e}{e-1} \frac{\beta}{(1-q\beta)n_m} \left(-1 + \frac{u'(w_{hm})^2}{u''(w_{hm})u(w_{hm})} \right)} \right) \\
& \Leftrightarrow \left(\frac{1}{u'(w_{lr})} - \frac{n_m}{u'(w_{lm})} \right) \frac{\partial U_r}{\partial \beta} = \frac{\theta q u(w_{hm})}{u'(w_{lm})} \left(\frac{1}{1 - \frac{\alpha_b e}{e-1} \frac{\beta}{(1-q\beta)n_m} \left(-1 + \frac{u'(w_{hm})^2}{u''(w_{hm})u(w_{hm})} \right)} \right).
\end{aligned} \tag{A.34}$$

The right-hand side of the expression above is positive, since $u''(w_{hm}) < 0$ and $u(w_{hm}) > 0$. From equation (A.28), it follows that $\frac{\partial U_r}{\partial \beta} < 0$.

If we instead have a partially tight labor market with $U_r = 0$, a change in β does not affect the utility of Red workers.

It then follows that: $\frac{\partial \pi(c_r^*)}{\partial \beta} = -\frac{\partial w_{lr}}{\partial \beta} - \frac{\partial w_{hr}}{\partial \beta} = -\frac{\partial U_r^*}{\partial \beta} \frac{1}{u'(w_{lr})} \geq 0$ for $\beta \geq \beta_{max}$. $\square$

References

- ARELLANO-BOVER, J., N. BIANCHI, S. LATTANZIO, AND M. PARADISI (2024): “One Cohort at a Time: A New Perspective on the Declining Gender Pay Gap,” NBER Working Paper 32612, National Bureau of Economic Research.
- ATHEY, S., C. AVERY, AND P. ZEMSKY (2000): “Mentoring and diversity,” *American Economic Review*, 90, 765–786.
- AURIOL, E., G. FRIEBEL, AND F. VON BIEBERSTEIN (2016): “The firm as the locus of social comparisons: Standard promotion practices versus up-or-out,” *Journal of Economic Behavior & Organization*, 121, 41–59.
- AURIOL, E. AND R. RENAULT (2008): “Status and incentives,” *The RAND Journal of Economics*, 39, 305–326.
- AXELSON, U. AND P. BOND (2015): “Wall Street Occupations,” *Journal of Finance*, 70, 1949–1996.
- AZMAT, G., V. CUÑAT, AND E. HENRY (2025): “Gender promotion gaps and career aspirations,” *Management Science*, 71, 2127–2141.
- BABCOCK, L., M. P. RECALDE, L. VESTERLUND, AND L. WEINGART (2017): “Gender Differences in Accepting and Receiving Requests for Tasks with Low Promotability,” *American Economic Review*, 107, 714–747.
- BAKER, G., M. GIBBS, AND B. HOLMSTROM (1994): “The Internal Economics of the Firm: Evidence from Personnel Data,” *Quarterly Journal of Economics*, 109, 881–919.

- BAKER, G. P., M. C. JENSEN, AND K. J. MURPHY (1988): “Compensation and incentives: Practice vs. theory,” *The Journal of Finance*, 43, 593–616.
- BARTOŠ, V., M. BAUER, J. CHYTILOVÁ, AND F. MATĚJKA (2016): “Attention discrimination: Theory and field experiments with monitoring information acquisition,” *American Economic Review*, 106, 1437–75.
- BECKER, G. S., K. M. MURPHY, AND I. WERNING (2005): “The equilibrium distribution of income and the market for status,” *Journal of Political Economy*, 113, 282–310.
- BENNEDSEN, M., E. SIMINTZI, M. TSOUTSOURA, AND D. WOLFENZON (2022): “Do Firms Respond to Gender Pay Gap Transparency?” *The Journal of Finance*, 77, 2051–2091.
- BENSON, A., D. LI, AND K. SHUE (2019): “Promotions and the Peter principle,” *The Quarterly Journal of Economics*, 134, 2085–2134.
- (2026): ““Potential” and the gender promotion gap,” *American Economic Review*, 116, 375–417.
- BERTRAND, M. AND K. F. HALLOCK (2001): “The gender gap in top corporate jobs,” *Industrial and Labor Relations Review*, 55, 3–21.
- BIRCAN, Ç., G. FRIEBEL, AND T. STAHL (2024): “Gender Promotion Gaps in Knowledge Work: The Role of Task Assignment in Teams,” EBRD Working Paper.
- BOHREN, J. A., A. IMAS, AND M. ROSENBERG (2019): “The dynamics of discrimination: Theory and evidence,” *American Economic Review*, 109, 3395–3436.
- BOSQUET, C., P.-P. COMBES, AND C. GARCÍA-PEÑALOSA (2019): “Gender and promotions: Evidence from academic economists in France,” *The Scandinavian Journal of Economics*, 121, 1020–1053.
- CALDWELL, S., I. HAEGELE, AND J. HEINING (2026): “Firm Pay, Amenities, and Inequality,” NBER Working Paper 35149, National Bureau of Economic Research, Cambridge, MA.
- CARD, D., A. R. CARDOSO, J. HEINING, AND P. KLINE (2018): “Firms and labor market inequality: Evidence and some theory,” *Journal of Labor Economics*, 36, S13–S70.
- CARD, D., A. R. CARDOSO, AND P. KLINE (2016): “Bargaining, sorting, and the gender wage gap: Quantifying the impact of firms on the relative pay of women,” *The Quarterly Journal of Economics*, 131, 633–686.
- CECCARELLI, M., C. HERPFER, AND S. ONGENA (2024): “Gender, performance, and promotion in the labor market for commercial bankers,” Swiss Finance Institute Research Paper.
- CORTÉS, P., J. HWANG, J. PAN, AND U. SCHÖNBERG (2026): “Gender Norms and the Labor Market,” NBER Working Paper 34716, National Bureau of Economic Research, Cambridge, MA.
- CULLEN, Z. AND R. PEREZ-TRUGLIA (2023): “The old boys’ club: Schmoozing and the gender gap,” *American Economic Review*, 113, 1703–1740.

- CZIRAKI, P. AND A. ROBERTSON (2021): “Credentials matter, but only for men: Evidence from the S&P 500,” Working paper, University of Toronto.
- DAVIES, S., E. VAN WESEP, AND B. WATERS (2024): “On the Magnification of Small Biases in Hiring,” *The Journal of Finance*, 79, 3623–3673.
- DOERINGER, P. B. AND M. J. PIORE (1971): *Internal Labor Markets and Manpower Analysis*, Lexington, MA: Heath Lexington Books.
- EWENS, M. AND X. GIROUD (2024): “Corporate Hierarchy,” Working Paper.
- FAHN, M. AND N. KLEIN (2025): “The Dynamics of Exploiting Overconfident Workers,” *Journal of Labor Economics*, forthcoming.
- FANG, H. AND A. MORO (2011): “Theories of statistical discrimination and affirmative action: A survey,” in *Handbook of Social Economics*, ed. by J. Benhabib, A. Bisin, and M. O. Jackson, North-Holland, vol. 1, chap. 5, 133–200.
- FERREIRA, D. AND R. NIKOLOVA (2024): “Prestige, promotion, and pay,” *Journal of Finance*, 79, 505–540.
- (2025): “Polarization, Purpose and Profit,” *Journal of Financial Economics*, 172, 104147.
- GIBBONS, R. AND M. WALDMAN (1999): “A Theory of Wage and Promotion Dynamics Inside Firms,” *Quarterly Journal of Economics*, 114, 1321–1358.
- GOLDIN, C. (2014): “A Pollution Theory of Discrimination: Male and Female Differences in Occupations and Earnings,” in *Human Capital in History: The American Record*, ed. by L. P. Boustan, C. Frydman, and R. A. Margo, Chicago: University of Chicago Press, 313–348.
- GOLDIN, C., S. P. KERR, C. OLIVETTI, AND E. BARTH (2017): “The Expanding Gender Earnings Gap: Evidence from the LEHD-2000 Census,” *American Economic Review: Papers and Proceedings*, 107, 110–114.
- HOFFMAN, M., L. B. KAHN, AND D. LI (2018): “Discretion in hiring,” *The Quarterly Journal of Economics*, 133, 765–800.
- HOPKINS, E. AND T. KORNIENKO (2004): “Running to Keep in the Same Place: Consumer Choice as a Game of Status,” *American Economic Review*, 94, 1085–1107.
- HOSPIDO, L., L. LAEVEN, AND A. LAMO (2019): “The gender promotion gap: Evidence from central banking,” *Review of Economics and Statistics*, 1–45.
- HUANG, R., E. J. MAYER, AND D. P. MILLER (2024): “Gender bias in promotions: Evidence from financial institutions,” *The Review of Financial Studies*, 37, 1685–1728.
- HUBER, K., V. LINDENTHAL, AND F. WALDINGER (2021): “Discrimination, managers, and firm performance: Evidence from “aryanizations” in Nazi Germany,” *Journal of Political Economy*, 129, 2455–2503.
- HUITFELDT, I., A. R. KOSTØL, J. NIMCZIK, AND A. WEBER (2023): “Internal Labor Markets: A Worker Flow Approach,” *Journal of Econometrics*, 233, 661–688.

- HUNEEUS, F., C. MILLER, C. NEILSON, AND S. D. ZIMMERMAN (2021): “Firm Sorting, College Major, and the Gender Earnings Gap,” *Working Paper, Princeton University*.
- KE, R., J. LI, AND M. POWELL (2018): “Managing Careers in Organizations,” *Journal of Labor Economics*, 36, 197–252.
- KLINE, P., E. K. ROSE, AND C. R. WALTERS (2022): “Systemic Discrimination Among Large U.S. Employers,” *The Quarterly Journal of Economics*, 137, 1963–2036.
- KREPS, D. M. AND J. A. SCHEINKMAN (1983): “Quantity precommitment and Bertrand competition yield Cournot outcomes,” *The Bell Journal of Economics*, 326–337.
- KUNZE, A. AND A. R. MILLER (2017): “Women helping women? Evidence from private sector data on workplace hierarchies,” *Review of Economics and Statistics*, 99, 769–775.
- LANG, K. AND J.-Y. K. LEHMANN (2012): “Racial Discrimination in the Labor Market: Theory and Empirics,” *Journal of Economic Literature*, 50, 959–1006.
- LANG, K., M. MANOVE, AND W. T. DICKENS (2005): “Racial discrimination in labor markets with posted wage offers,” *American Economic Review*, 95, 1327–1340.
- LAZEAR, E., K. SHAW, AND C. STANTON (2018): “Who gets hired? The importance of competition among applicants,” *Journal of Labor Economics*, 36, S133–S181.
- LAZEAR, E. P. AND S. ROSEN (1981): “Rank-order tournaments as optimum labor contracts,” *Journal of Political Economy*, 89, 841–864.
- (1990): “Male-female wage differentials in job ladders,” *Journal of Labor Economics*, 8, S106–S123.
- MAS, A. (2025): “Non-Wage Amenities,” *Handbook of Labor Economics*, 6.
- MATSA, D. A. AND A. R. MILLER (2011): “Chipping away at the glass ceiling: Gender spillovers in corporate leadership,” *American Economic Review*, 101, 635–639.
- MINNI, V., K.-T. NGUYEN, AND H. SARSONS (2025): “Managers and the Cultural Transmission of Gender Norms,” *Working Paper*.
- MORCHIO, I. AND C. MOSER (2026): “The Gender Pay Gap: Micro Sources and Macro Consequences,” *The American Economic Review*, 116, 1765—1810.
- ONUCHIC, P. (2023): “Recent contributions to theories of discrimination,” *Working paper, Oxford University*.
- ONUCHIC, P. AND D. RAY (2023): “Signaling and discrimination in collaborative projects,” *American Economic Review*, 113, 210–252.
- PIKULINA, E. AND D. FERREIRA (2026): “Subtle Discrimination,” *Journal of Finance*, 81, 329–369.
- PRENDERGAST, C. (1993): “The role of promotion in inducing specific human capital acquisition,” *The Quarterly Journal of Economics*, 108, 523–534.

- RAY, D. AND A. ROBSON (2012): “Status, Intertemporal Choice, and Risk-Taking,” *Econometrica*, 80, 1505–1531.
- RONCHI, M. AND N. SMITH (2026): “Daddy’s girl: Daughters, managerial decisions, and gender inequality,” *Review of Economic Studies*, forthcoming.
- ROSEN, S. (1986): “The theory of equalizing differences,” *Handbook of Labor Economics*, 1, 641–692.
- SINISCALCHI, M. AND P. VERONESI (2021): “Self-image bias and lost talent,” Working paper, NBER.
- WALDMAN, M. (1984): “Job assignments, signaling, and efficiency,” *The RAND Journal of Economics*, 15, 255–267.
- (2013): “Classic promotion tournaments versus market-based tournaments,” *International Journal of Industrial Organization*, 31, 198–210.
- WRIGHT, R., P. KIRCHER, B. JULIEN, AND V. GUERRIERI (2021): “Directed search and competitive search equilibrium: A guided tour,” *Journal of Economic Literature*, 59, 90–148.

Internet Appendix for “Biased Promotions”

IA.I Characterizing the equilibrium contracts

A fully or partially tight labor market equilibrium is a set of two unique contracts $\{c_r^*, c_m^*\}$ and a belief function q , such that firm $\tau_m \in [0, s^*F]$ offers contract $c_m^* = (n_m^*, w_{lm}^*, w_{hm}^*)$, firm $\tau_r \in (s^*F, F]$ offers contract $c_r^* = (1, w_{lr}^*, w_{hr}^*)$, and $s^* \in (0, 1]$.

In a fully tight labor market, Mixed firms attract Blue and Red workers. In the proof of Proposition 4, we show that in Mixed firms, only the participation constraint of the Red workers binds.

The equilibrium must then satisfy seven conditions:

- i. Workers' lifetime utilities are

$$\begin{aligned} U_r^* &:= u(w_{lm}^*) + \frac{1-q^*\beta}{n_m^*} \theta u(w_{hm}^*) \\ U_b^* &:= u(w_{lm}^*) + \frac{1+(1-q^*)\beta}{n_m^*} \theta u(w_{hm}^*) \end{aligned} \tag{IA.1}$$

where $q^* := q(c_m^*, C_{-\tau_m}^*)$.

- ii. Type- m firms choose c_m^* optimally:

$$\pi(c_m^*) = \max_{n_m, w_{lm}, w_{hm}} R_h - w_{hm} + n_m(R_l - w_{lm}) \tag{IA.2}$$

subject to

$$u(w_{lm}) + \frac{1-q^*\beta}{n_m} \theta u(w_{hm}) \geq U_r^*. \tag{IA.3}$$

- iii. Type- r firms choose c_r^* optimally:

$$\pi(c_r^*) := \max_{w_{lr}, w_{hr}} R_h - w_{hr} + R_l - w_{lr} \tag{IA.4}$$

subject to

$$u(w_{lr}) + \theta u(w_{hr}) \geq U_r^*. \quad (\text{IA.5})$$

iv. Firms' choice of contract type must be optimal:

$$\begin{cases} \text{If } \pi(c_m^*) > \pi(c_r^*), \text{ then } s^* = 1 \\ \text{If } \pi(c_m^*) = \pi(c_r^*), \text{ then } s^* < 1 \end{cases} \quad (\text{IA.6})$$

v. If $s^* < 1$, Red workers are indifferent between contracts c_m^* and c_r^* :

$$U_r^* = u(w_{lr}^*) + \theta u(w_{hr}^*). \quad (\text{IA.7})$$

vi. Labor markets clear:

$$\begin{cases} s^* q^* n_m^* F = \alpha_b E \\ (1 - s^* + s^*(1 - q^*) n_m^*) F = (1 - \alpha_b) E \end{cases} . \quad (\text{IA.8})$$

vii. There is no profitable deviation: $\nexists c_\tau^d$ such that $\mu_{b\tau}(c_\tau^d, C_{-\tau}^*) = n_\tau^d q(c_\tau^d, C_{-\tau}^*)$ and $\pi(c_\tau^d) > \pi(c_m^*)$.

Since in Mixed firms the participation constraint of Red binds, a deviation contract $c^d \notin \{c_m^*, c_r^*\}$ that meets the Red constraints is sub-optimal. Thus, even if there are rational beliefs that sustain such deviations, they are not profitable. Suppose now that c^d violates both participation constraints at belief q^* . In this case, no one will be assigned to this contract, implying that any q is rational for this deviation. Thus, we set $q(c) = q^*$ for any contract that attracts no worker. It follows that the profit is lower in this deviation than in the equilibrium contract.

The only remaining case is a deviation that, under q^* , violates the Red constraint but not the Blue constraint. In this case, the firm offers a risky career path ($n^d > 1$) only to Blue workers. Under this deviation, the only rational belief is $q = 1$. Note that if a

contract violates the Red constraint at q^* , it will also violate it at $q = 1$. It then follows that Condition (vii.) holds if and only if

$$\pi(c_m^*) \geq \max_c \pi(c) \text{ s.t. } U_b(c; 1) \geq U_b^*, \quad (\text{IA.9})$$

where $U_b(c; q)$ is b 's utility from contract c and belief q . In words, the absence of a profitable deviation requires the equilibrium profit to be greater than the maximum profit subject to hiring only Blue workers.

In a partially tight labor market only Blue workers apply to Mixed firms. Below we focus on the modified conditions from the list (i)-(vii) presented before:

i. Workers' lifetime utilities are

$$\begin{aligned} U_r^* &:= u(w_{lr}^*) + \theta u(w_{hr}^*) \\ U_b^* &:= u(w_{lm}^*) + \frac{1}{n_m^*} \theta u(w_{hm}^*) \end{aligned} \quad (\text{IA.10})$$

that is $q^* := 1$ and $U_r^* = 0$.

ii. Type- m firms choose c_m^* optimally:

$$\pi(c_m^*) = \max_{n_m, w_{lm}, w_{hm}} R_h - w_{hm} + n_m(R_l - w_{lm}) \quad (\text{IA.11})$$

subject to

$$u(w_{lm}^*) + \frac{1}{n_m^*} \theta u(w_{hm}^*) \geq U_b^*. \quad (\text{IA.12})$$

Therefore, the optimal contract must satisfy the following marginal conditions:

$$R_l = w_{lm}^* + \frac{U_b^* - u(w_{lm}^*)}{u'(w_{lm}^*)} \quad (\text{IA.13})$$

$$u'(w_{lm}^*) = \theta u'(w_{hm}^*) \quad (\text{IA.14})$$

Notice that the wages are independent of β .

Conditions (iii.) and (iv.) are as before.

v. Red workers do not deviate and apply to Mixed firms:

$$u(w_{lm}^*) + \frac{(1 - \beta)\theta u(w_{hm}^*)}{n_m^*} < 0 \quad (\text{IA.15})$$

vi. Labor markets clear:

$$\begin{cases} s^* n_m^* F = \alpha_b E \\ (1 - s^*) F \leq (1 - \alpha_b) E \end{cases} \quad (\text{IA.16})$$

We can further rewrite the non-deviation condition, using the marginal conditions for the optimal contract c_m^* and the expression for U_b^* :

$$u(w_{lm}^*) + \frac{(1 - \beta)\theta u(w_{hm}^*)}{\frac{\theta u(w_{hm}^*)}{u'(w_{lm}^*)(R_l - w_{lm}^*)}} < 0 \Leftrightarrow u(w_{lm}^*) + (1 - \beta)u'(w_{lm}^*)(R_l - w_{lm}^*) < 0 \quad (\text{IA.17})$$

Since the wages w_{lm}^* is independent of β , the left-hand side of the non deviation condition is decreasing in β . For $\beta = 1$, non-deviation requires $u(w_{lm}^*) < 0$. From Assumption 1(ii), $u(R_l) < 0$ and since $w_{lm}^* < R_l$ and $u'(\cdot) > 0$, it follows that $u(w_{lm}^*) < 0$ and there exists a $\tilde{\beta} < 1$, where $\tilde{\beta}$ is given by $u(w_{lm}^*) + (1 - \tilde{\beta})u'(w_{lm}^*)(R_l - w_{lm}^*) = 0$, such that for $\beta > \tilde{\beta}$ only Blue workers apply to contract c_m^* .

IA.II Equilibrium for $\beta < \beta_{min}$

For completeness, in this section, we present the comparative static with respect to β for $\beta < \beta_{min}$. From Proposition 5, if $\beta < \beta_{min}$ we know that $s^* = 1$, that is all firms are Mixed firms and hire both Blue and Red workers. The span of control in all firms is the same and equal to $n_m^* = \frac{E}{F}$ and the fraction of blue workers employed by each firm is the same as the fraction of blue workers in the worker population, i.e., $q^* = \alpha_b$. We now discuss the effect of the bias on wages, utilities, and profits.

Figure IA.1 is the equivalent of Figure 2B, including the wages in Mixed firms for values of $\beta \in [0, \beta_{min}]$. An increase in the bias β decreases the outside option of Red workers U_r^* . A lower U_r^* leads to a higher entry-level wage w_{lm}^* . Differently from the case with two types of firms, as β increases, the span of control (n_m^*) remains constant and the wages in the top job w_{hm}^* decrease.

Figure IA.2 shows that, as expected, the equilibrium utility of Red workers decreases with the level of the bias β . The effect of the bias on the utility of Blue is more subtle. On the one hand, higher β increases the probability of promotion, on the other hand, it decreases the lifetime utility of Red workers thus leading to contracts that are less favourable to all workers. In the case where all firms are Mixed, the direct effect is stronger and thus leads to an overall increase in the lifetime utility of Blue workers.

Even though β negatively affects the bargaining power of Red workers, the distortion that the bias creates relative to the efficient contract dominates and for low values of β the firms profits decrease, as shown in Figure IA.3.




Figure IA.1: Wages in Mixed firms, w_{hm}^* and w_{lm}^* .

Note: This figure presents the equilibrium wages in Mixed firms for top job h , w_{hm}^* (solid line), and entry-level jobs l , w_{lm}^* (dashed line), as functions of subtle bias, $\beta \geq 0$, for $\theta = 10.0$, $R_l = 0.75$, $\alpha_b = 0.5$, and $e = 2.0$.




Figure IA.2: Utilities of Blue and Red workers, U_b^* and U_r^* .

Note: This figure presents the equilibrium utilities of Blue workers (solid blue line) and Red workers (dashed red line), U_b^* and U_r^* , as functions of subtle bias, $\beta \geq 0$, for $\theta = 10.0$, $R_l = 0.75$, $\alpha_b = 0.5$, and $e = 2.0$. The horizontal dotted line represents the equilibrium utility, $U_b^* = U_r^* = U^*$, in the benchmark case, $\beta = 0$.




Figure IA.3: Firm profit, π^* .

Note: This figure presents the equilibrium profit of Mixed firms, π^* , as a function of subtle bias, $\beta \geq 0$, for $\theta = 10.0$, $R_l = 0.75$, $R_h = 4.0$, $\alpha_b = 0.5$, and $e = 2.0$. The horizontal dotted lines represent the equilibrium profit, π^* , in the benchmark case, $\beta = 0$.

IA.III Bias-reducing technology

In this section, we consider an extension of the model, where the mass of firms F is given but firms have the possibility to acquire a bias-reducing technology or alternatively take a costly action that allows them to credibly reveal to workers that $\beta = 0$, before choosing a contract c_τ . The per-period cost of the technology is I . A firm τ chooses to acquire the technology only if:

$$\pi_0(c_0^*) - I \geq \pi_\beta(c_\beta^*), \quad (\text{IA.18})$$

where c_0^* and c_β^* denote the firm's optimal contracts when it is perceived to have no promotion bias and bias β , respectively, and $\pi_0(c_0^*)$ and $\pi_\beta(c_\beta^*)$ are the corresponding profits.

Note that in firms with $\beta = 0$, for a given contract, Blue and Red workers have the same expected utility. In biased firms Blue workers have strictly higher utility than Red workers for the same contract.

In this section, we are asking whether allowing firms to acquire a costly bias-reducing technology could eliminate the equilibrium effect of the subtle bias and restore efficiency. The next proposition shows that it cannot.

Proposition IA.1. *For any $I > 0$, there is no equilibrium in which all Red workers are employed by firms with $\beta = 0$ and the equilibrium contract is c^* .*

Proof. If all Red work in firms with $\beta = 0$, then in firms with $\beta > 0$ there are only Blue workers. With only Blue workers in biased firms, the bias does not affect the expected utility of Blue workers and therefore their contract. In order to eliminate the effect of the bias it must be that $U_r^* = U_b^*$. In that case the contract offered by both types of firms is c^* (as defined in Proposition 2), so that $c_0^* = c_\beta^* = c^*$. This implies that $\pi_0(c_0^*) = \pi_\beta(c_\beta^*)$, and therefore condition (IA.18) is violated for any $I > 0$. $\square$

Proposition IA.1 implies that a costly bias-reducing technology, even if available, cannot fully eliminate the promotion bias. The technology is only valuable when the profits of biased firms are distorted by the bias. If all Red work in firms offering fair promotion

lotteries, the bias plays no role in the firms with $\beta > 0$ thus eliminating the advantage from acquiring the technology.

About ECGI

ECGI is an international non-profit membership organisation that provides a platform for debate and dialogue among academics, policymakers, and business leaders, with a focus on major corporate governance, ESG, and stewardship issues.

The views expressed in this working paper are those of the authors, not those of the ECGI or its members.

www.ecgi.global

ECGI Working Paper Series in Finance

Editorial Board

Editor

Nadya Malenko, Professor of Finance, Boston College

Consulting Editors

Renée Adams, Professor of Finance, University of Oxford

Franklin Allen, Professor of Finance and Economics, Imperial London

Julian Franks, Professor of Finance, London Business School

Mireia Giné, Professor of Finance, IESE Business School

Marco Pagano, Professor of Economics, Facoltà di Economia Università di Napoli
Federico II

Series Manager

Asif Malik, ECGI Working Paper Series Manager

<https://www.ecgi.global/publications/working-papers>